

The astrometric core solution for the Gaia mission

Overview of models, algorithms and software implementation

L. Lindegren¹, U. Lammers², D. Hobbs¹, W. O’Mullane², U. Bastian³, and J. Hernández²

¹ Lund Observatory, Lund University, Box 43, SE-22100 Lund, Sweden
e-mail: Lennart.Lindegren, David.Hobbs@astro.lu.se

² European Space Agency (ESA), European Space Astronomy Centre (ESAC), P.O. Box (Apdo. de Correos) 78, ES-28691 Villanueva de la Cañada, Madrid, Spain
e-mail: Uwe.Lammers, William.O’Mullane, Jose.Hernandez@sciops.esa.int

³ Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg, Mönchhofstr. 12–14, DE-69120 Heidelberg, Germany
e-mail: bastian@ari.uni-heidelberg.de

Received 17 August 2011 / Accepted 25 November 2011

ABSTRACT

Context. The Gaia satellite will observe about one billion stars and other point-like sources. The astrometric core solution will determine the astrometric parameters (position, parallax, and proper motion) for a subset of these sources, using a global solution approach which must also include a large number of parameters for the satellite attitude and optical instrument. The accurate and efficient implementation of this solution is an extremely demanding task, but crucial for the outcome of the mission.

Aims. We aim to provide a comprehensive overview of the mathematical and physical models applicable to this solution, as well as its numerical and algorithmic framework.

Methods. The astrometric core solution is a simultaneous least-squares estimation of about half a billion parameters, including the astrometric parameters for some 100 million well-behaved so-called primary sources. The global nature of the solution requires an iterative approach, which can be broken down into a small number of distinct processing blocks (source, attitude, calibration and global updating) and auxiliary processes (including the frame rotator and selection of primary sources). We describe each of these processes in some detail, formulate the underlying models, from which the observation equations are derived, and outline the adopted numerical solution methods with due consideration of robustness and the structure of the resulting system of equations. Appendices provide brief introductions to some important mathematical tools (quaternions and B-splines for the attitude representation, and a modified Cholesky algorithm for positive semidefinite problems) and discuss some complications expected in the real mission data.

Results. A complete software system called AGIS (Astrometric Global Iterative Solution) is being built according to the methods described in the paper. Based on simulated data for 2 million primary sources we present some initial results, demonstrating the basic mathematical and numerical validity of the approach and, by a reasonable extrapolation, its practical feasibility in terms of data management and computations for the real mission.

Key words. Astrometry – Methods: data analysis – Methods: numerical – Space vehicles: instruments

1. Introduction

The space astrometry mission Gaia, planned to be launched by the European Space Agency (ESA) in 2013, will determine accurate astrometric data for about one billion objects in the magnitude range from 6 to 20. Accuracies of 8–25 micro-arcsec (μ as) are typically expected for the trigonometric parallaxes, positions at mean epoch, and annual proper motions of simple (i.e., apparently single) stars down to 15th magnitude. The astrometric data are complemented by photometric and spectroscopic information collected with dedicated instruments on board the Gaia satellite. The mission will result in an astronomical database of unprecedented scope, accuracy and completeness becoming available to the scientific community around 2021.

The original interferometric concept for a successor mission to Hipparcos, called GAIA (Global Astrometric Interferometer for Astrophysics), was described by Lindegren & Perryman (1996) but has since evolved considerably by the incorporation of novel ideas (Høg 2008) and as a result of industrial studies conducted under ESA contracts (Perryman et al. 2001). The mis-

sion, now in the final integration phase with EADS Astrium as prime contractor, is no longer an interferometer but has retained the name Gaia, which is thus no acronym. For some brief but up-to-date overviews of the mission, see Lindegren et al. (2008) and Lindegren (2010). The scientific case is most comprehensively described in the proceedings of the conference The Three-Dimensional Universe with Gaia (Turon et al. 2005).

In parallel with the industrial development of the satellite, the Gaia Data Processing and Analysis Consortium (DPAC; Mignard et al. 2008) is charged with the task of developing and running a complete data processing system for analysing the satellite data and constructing the resulting database (‘Gaia Catalogue’). This task is extremely difficult due to the large quantities of data involved, the complex relationships between different kinds of data (astrometric, photometric, spectroscopic) as well as between data collected at different epochs, the need for complex yet efficient software systems, and the interaction and sustained support of many individuals and groups over an extended period of time.

A fundamental part of the data processing task is the astrometric core solution, currently under development in DPAC's Coordination Unit 3 (CU3), 'Core Processing'. Mathematically, the astrometric core solution is a simultaneous determination of a very large number of unknowns representing three kinds of information: (i) the astrometric parameters for a subset of the observed stars, representing the astrometric reference frame; (ii) the instrument attitude, representing the accurate celestial pointing of the instrument axes in that reference frame as a function of time; and (iii) the geometric instrument calibration, representing the mapping from pixels on the CCD detectors to angular directions relative to the instrument axes. Although the astrometric core solution is only made for a subset of the stars, the resulting celestial reference frame, attitude and instrument calibration are fundamental inputs for the processing of all observations. Optionally, a fourth kind of unknowns, the global parameters, may be introduced to describe for example a hypothetical deviation from General Relativity.

We use the term 'source' to denote any astronomical object that Gaia detects and observes as a separate entity. The vast majority of the Gaia sources are ordinary stars, many of them close binaries or the components of wide systems, but some are non-stellar (for example asteroids and quasars). Nearly everywhere in this paper, one can substitute 'star' for 'source' without distortion; however, for consistency with established practice in the Gaia community we use 'source' throughout.

While the total number of distinct sources that will be observed by Gaia is estimated to slightly more than one billion, only a subset of them shall be used in the astrometric core solution. This subset, known as the 'primary sources', is selected to be astrometrically well-behaved (see Sect. 6.2) and consists of (effectively) single stars and extragalactic sources (quasars and AGNs) that are sufficiently stable and point-like. We assume here that the number of primary sources is about 10^8 , i.e., roughly one tenth of the total number of objects, although in the end it is possible that an even larger number will be used.

In comparison with many other parts of the Gaia data processing, the astrometric core solution is in principle simple, mainly because it only uses a subset of the observations (namely, those of the primary sources), which can be accurately modelled in a relatively straightforward way. In practice, the problem is however formidable: the total number of unknowns is of the order of 5×10^8 , the solution uses some 10^{11} individual observations extracted from some 70 Terabyte (TB) of raw satellite data, and the entangled nature of the data excludes a direct solution. A feasible approach has nevertheless been found, including a precise mathematical formulation, practical solution method, and efficient software implementation. It is the aim of this paper to provide a comprehensive overview of this approach.

Concerning notations we have followed the usual convention to denote all non-scalar entities (vectors, tensors, matrices, quaternions) by boldfaced characters. Lower-case bold italics (\mathbf{a}) are used for vectors and one-dimensional column matrices; upper-case bold italics (\mathbf{A}) usually denote two-dimensional matrices. Following Murray (1983) the prime (') signifies both matrix transpose (\mathbf{a}' , \mathbf{A}') and scalar multiplication of vectors; thus $\|\mathbf{a}\| = (\mathbf{a}'\mathbf{a})^{1/2}$ defines the magnitude of the vector \mathbf{a} as well as the Euclidean norm of the column matrix \mathbf{a} . Angular brackets denote normalization to unit length, as in $\langle \mathbf{x} \rangle = \mathbf{x}/\|\mathbf{x}\|^{-1}$. In this notation, no special distinction is thus made between vectors as physical entities (also known as geometric, spatial or Euclidean vectors) on one hand, and their numerical representations in some coordinate system as column matrices (also known as list vectors) on the other hand. Moreover, list vectors can of course have any

dimension: $\mathbf{a} \in \mathbb{R}^n$. In the coordinate system whose axes are aligned with unit vectors \mathbf{x} , \mathbf{y} , and \mathbf{z} , the components of the arbitrary vector \mathbf{a} are given by $a_x = \mathbf{x}'\mathbf{a}$, $a_y = \mathbf{y}'\mathbf{a}$, and $a_z = \mathbf{z}'\mathbf{a}$; if the coordinate system is represented by the vector triad $\mathbf{S} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$, these components are given by the column matrix $\mathbf{S}'\mathbf{a}$ (cf. Appendix A in Murray 1983). This notation, although perhaps unfamiliar to many readers, provides a convenient and unambiguous framework for representing and transforming spatial vectors in different coordinate systems. For a vector-valued function $\mathbf{f}(\mathbf{x})$, $\partial\mathbf{f}/\partial\mathbf{x}'$ denotes the $\dim(\mathbf{f}) \times \dim(\mathbf{x})$ matrix whose (i, j) th element is $\partial f_i / \partial x_j$. Quaternions follow their own algebra (see Appendix A for a brief introduction) and must not be confused with vectors/matrices; quaternions are therefore denoted by bold Roman characters (\mathbf{a}). When taking a derivative with respect to the quaternion \mathbf{a} , $\partial\mathbf{x}/\partial\mathbf{a}$ denotes the 4×1 matrix of derivatives with respect to the quaternion components; $\partial\mathbf{x}/\partial\mathbf{a}'$ is the transposed matrix. Tables of acronyms and variables are provided in Appendix E.

2. Outline of the approach

The astrometric principles for Gaia were outlined already in the Hipparcos Catalogue (ESA 1997, Vol. 3, Ch. 23) where, based on the accumulated experience of the Hipparcos mission, a view was cast to the future. The general principle of a global astrometric data analysis was succinctly formulated as the minimization problem:

$$\min_{\mathbf{s}, \mathbf{n}} \|\mathbf{f}^{\text{obs}} - \mathbf{f}^{\text{calc}}(\mathbf{s}, \mathbf{n})\|_{\mathcal{M}}, \quad (1)$$

with a slight change of notation to adapt to the present paper. Here \mathbf{s} is the vector of unknowns (parameters) describing the barycentric motions of the ensemble of sources used in the astrometric solution, and \mathbf{n} is a vector of 'nuisance parameters' describing the instrument and other incidental factors which are not of direct interest for the astronomical problem but are nevertheless required for realistic modelling of the data. The observations are represented by the vector \mathbf{f}^{obs} which could for example contain the measured detector coordinates of all the stellar images at specific times. $\mathbf{f}^{\text{calc}}(\mathbf{s}, \mathbf{n})$ is the observation model, e.g., the expected detector coordinates calculated as functions of the astrometric and nuisance parameters. The norm is calculated in a metric \mathcal{M} defined by the statistics of the data; in practice the minimization will correspond to a weighted least-squares solution with due consideration of robustness issues (see Sect. 3.6).

While Eq. (1) encapsulates the global approach to the data analysis problem, its actual implementation will of course depend strongly on the specific characteristics of the Gaia instrument and mission as well as on the practical constraints on the data processing task.

At the heart of the problem is the modelling of data represented by the function $\mathbf{f}^{\text{calc}}(\mathbf{s}, \mathbf{n})$ in Eq. (1). This is schematically represented in the shaded part of the diagram in Fig. 1. It shows the main steps for calculating the expected CCD output in terms of the various parameters. The data processing, effecting the minimization in Eq. (1), is represented in the right part of the diagram. In subsequent sections we describe in some detail the main elements depicted in Fig. 1. The astrometric (source) parameters are represented by the vector \mathbf{s} , while the nuisance parameters are of three kinds: the attitude parameters \mathbf{a} , the geometric calibration parameters \mathbf{c} , and the global parameters \mathbf{g} . The observables \mathbf{f} are represented by the pixel coordinates (κ, μ) of point source images on Gaia's CCDs. (In the actual implementation of the approach, the minimization problem is formu-

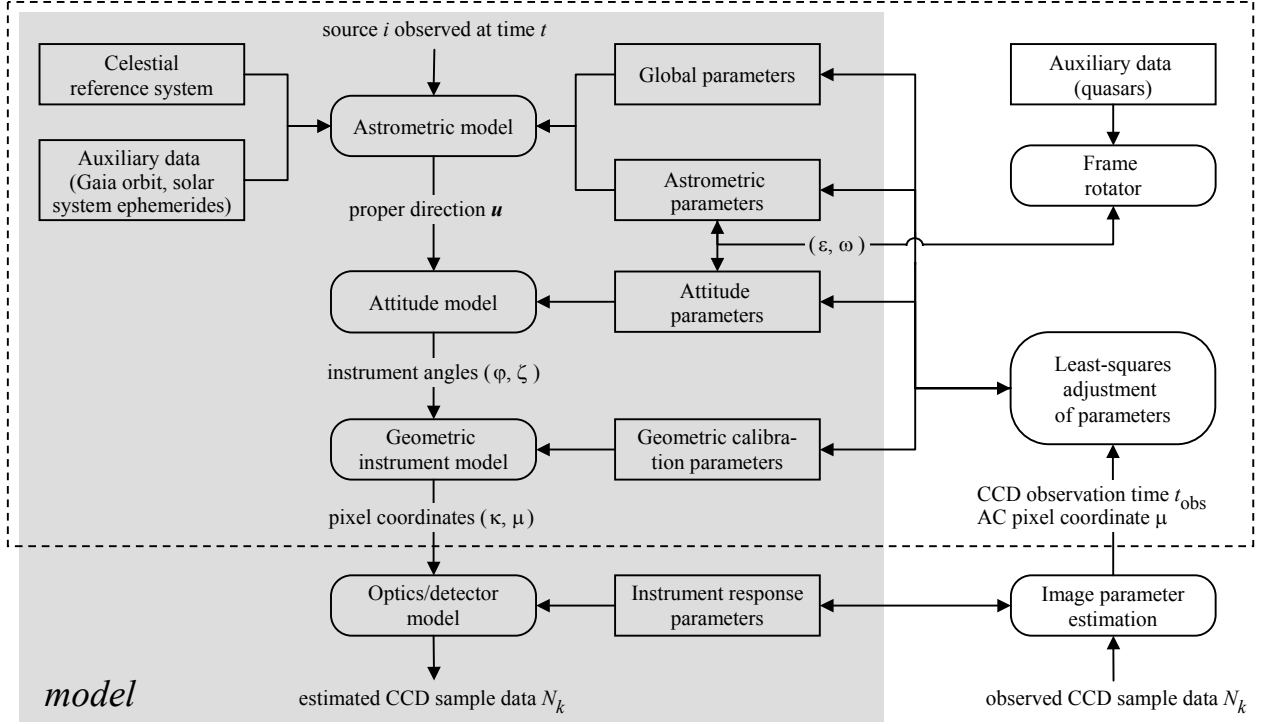


Fig. 1. Schematic representation of the main elements of the astrometric data analysis. In the shaded area is the mathematical model that allows to calculate the position of the stellar image in pixel coordinates, and hence the expected CCD data, for any given set of model parameters. To the right are the processes that fit this model to the observed CCD data by adjusting the parameters in the rectangular boxes along the middle. This paper is primarily concerned with the geometrical part of the analysis contained in the dashed box. However, a brief outline of the CCD data modelling and processing (bottom part of the diagram) is given in Sect. 3.5 and Appendix D.2.

lated in terms of the field angles η , ζ rather than in the directly measured pixel coordinates κ , μ ; see Sects. 3.3 and 3.4.)

The various elements of the astrometric solution are described in some detail in the subsequent sections. Section 3 provides the mathematical framework needed to model the Gaia observations and setting up the least-squares equations for the astrometric solution. In the interest of clarity and overview we omit in this description certain complications that need to be considered in the final data processing system; these are instead briefly discussed in Appendix D. In Sect. 4 we describe the iterative solution method in general terms, and then provide, in Sects. 5 and 6, the mathematical details of the most important procedures. Section 7 outlines the existing implementation of the solution and presents the results of a demonstration run based on simulated observations of about 2 million primary sources. Appendices A to C provide brief introductions to three mathematical tools that are particularly important for the subsequent development, namely the use of quaternions for representing the instantaneous satellite attitude, the B-spline formalism used to model the attitude as a function of time, and a modified Cholesky algorithm for the decomposition of positive semidefinite normal matrices.

3. Mathematical formulation of the basic observation model

3.1. Reference systems

The high astrometric accuracy aimed for with Gaia makes it necessary to use general relativity for modelling the data. This implies a precise and consistent formulation of the different reference systems used to describe the motion of the observer (Gaia), the motion of the observed object (source), the propagation of light from the source to the observer, and the transformations between these systems. The formulation adopted for Gaia (Klioner 2003, 2004) is based on the parametrized post-Newtonian (PPN) version of the relativistic framework adopted in 2000 by the International Astronomical Union (IAU); see Soffel et al. (2003). In this section only some key concepts from this formulation are introduced.

The orbit of Gaia and the light propagation from the source to Gaia are entirely modelled in the Barycentric Celestial Reference System (BCRS) whose spatial axes are aligned with the International Celestial Reference System (ICRS, Feissel & Mignard 1998). The associated time coordinate is the barycentric coordinate time (TCB). Throughout this paper, all time variables denoted t (with various subscripts) must be interpreted as TCB. The ephemerides of solar-system bodies (including the Sun and the Earth) are also expressed in this reference system. Even the motions of the stars and extragalactic objects are formally modelled in this system, although for practical reasons certain effects of the light-propagation time are conventionally ignored in this model (Sect. 3.2).

In order to model the rotation (attitude) of Gaia and the celestial direction of the light rays as observed by Gaia, it is expedient to introduce also a co-moving celestial reference system having its origin at the centre of mass of the satellite and a coordinate time equal to the proper time at Gaia. This is known as the Centre-of-Mass Reference System (CoMRS) and the associated time coordinate is Gaia Time (T_G). Klioner (2004) demonstrates how the CoMRS can be constructed in the IAU 2000 framework in exact analogy with the Geocentric Celestial Reference System (GCRS), only for a massless particle (Gaia) instead of the Earth. Like the BCRS and GCRS, the CoMRS is kinematically non-rotating, and its spatial axes are aligned with the ICRS. The celestial reference system (either the CoMRS or the ICRS depending on whether the origin is at Gaia or the solar-system barycentre) will in the following be represented by the vector triad $\mathbf{C} = [\mathbf{X} \ \mathbf{Y} \ \mathbf{Z}]$, where \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are orthogonal unit vectors aligned with the axes of the celestial reference system. That is, \mathbf{X} points towards the origin ($\alpha = \delta = 0$), \mathbf{Z} towards the north celestial pole ($\delta = +90^\circ$), and $\mathbf{Y} = \mathbf{Z} \times \mathbf{X}$ points to ($\alpha = 90^\circ, \delta = 0$).

In the CoMRS the attitude of the satellite is a spatial rotation in three dimensions, and can therefore be described purely classically, for example using quaternions (Sect. 3.3 and Appendix A). The rotated reference system, aligned with the instrument axes, is known as the Scanning Reference System (SRS). Its spatial \mathbf{x} , \mathbf{y} , \mathbf{z} axes (Fig. 2) are defined in terms of the two viewing directions of Gaia \mathbf{f}_P (in the centre of the preceding field of view, PFOV) and \mathbf{f}_F (in the centre of the following field of view, FFOV) as

$$\mathbf{x} = \langle \mathbf{f}_F + \mathbf{f}_P \rangle, \quad \mathbf{z} = \langle \mathbf{f}_F \times \mathbf{f}_P \rangle, \quad \mathbf{y} = \mathbf{z} \times \mathbf{x}. \quad (2)$$

(The precise definition of \mathbf{f}_P and \mathbf{f}_F is implicit in the geometric instrument model; see Sect. 3.4.) The SRS is represented by the vector triad $\mathbf{S} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$.

For determining the orbit of Gaia and calibrating the on-board clock, it is also necessary to model the radio ranging and other ground-based observations of the Gaia spacecraft in the same relativistic framework. For this, we also need the GCRS. These aspects of the Gaia data processing are, however, not discussed in this paper.

3.2. Astrometric model

The astrometric model is a recipe for calculating the proper direction $\mathbf{u}_i(t)$ to a source (i) at an arbitrary time of observation (t) in terms of its astrometric parameters s_i and various auxiliary data, assumed to be known with sufficient accuracy. The auxiliary data include an accurate barycentric ephemeris of the Gaia satellite, $\mathbf{b}_G(t)$, including its coordinate velocity $d\mathbf{b}_G/dt$, and ephemerides of all other relevant solar-system bodies.

For the astrometric core solution every source is assumed to move with uniform space velocity relative to the solar-system barycentre. In the BCRS its spatial coordinates at time t_* are therefore given by

$$\mathbf{b}_i(t_*) = \mathbf{b}_i(t_{\text{ep}}) + (t_* - t_{\text{ep}})\mathbf{v}_i \quad (3)$$

where t_{ep} is an arbitrary reference epoch and $\mathbf{b}_i(t_{\text{ep}})$, \mathbf{v}_i define six kinematic parameters for the motion of the source. The six astrometric parameters in s_i are merely a transformation of the kinematic parameters into an equivalent set better suited for the analysis of the observations. The six astrometric parameters are:

α_i the barycentric right ascension at the adopted reference epoch t_{ep}

δ_i the barycentric declination at epoch t_{ep}

ϖ_i the annual parallax at epoch t_{ep}

$\mu_{\alpha^*i} = (d\alpha_i/dt) \cos \delta_i$ the proper motion in right ascension at epoch t_{ep}

$\mu_{\delta i} = d\delta_i/dt$ the proper motion in declination at epoch t_{ep}

$\mu_{ri} = v_{ri}\varpi_i/A_u$ the ‘radial proper motion’ at epoch t_{ep} , i.e., the radial velocity of the source (v_{ri}) expressed in the same unit as the transverse components of proper motion (A_u = astronomical unit).

As explained in Sect. 3.1, the astrometric parameters refer to the ICRS and the time coordinate used is TCB. The reference epoch t_{ep} is preferably chosen to be near the mid-time of the mission in order to minimize statistical correlations between the position and proper motion parameters.

The transformation between the kinematic and the astrometric parameters is non-trivial (Klioner 2003), mainly as a consequence of the practical need to neglect most of the light-propagation time $t - t_*$ between the emission of the light at the source (t_*) and its reception at Gaia (t). This interval is typically many years and its value, and rate of change (which depends on the radial velocity of the source), will in general not be known with sufficient accuracy to allow modelling of the motion of the source directly in terms of its kinematic parameters according to Eq. (3). The proper motion components μ_{α^*i} , $\mu_{\delta i}$ and radial velocity v_{ri} correspond to the ‘apparent’ quantities discussed by Klioner (2003, Sect. 8).

The coordinate direction to the source at time t is calculated with the same ‘standard model’ as was used for the reduction of the Hipparcos observations (ESA 1997, Vol. 1, Sect. 1.2.8), namely

$$\bar{\mathbf{u}}_i(t) = \langle \mathbf{r}_i + (t_B - t_{\text{ep}})(\mathbf{p}_i\mu_{\alpha^*i} + \mathbf{q}_i\mu_{\delta i} + \mathbf{r}_i\mu_{ri}) - \varpi_i\mathbf{b}_G(t)/A_u \rangle \quad (4)$$

where the angular brackets signify vector length normalization, and $[\mathbf{p}_i \ \mathbf{q}_i \ \mathbf{r}_i]$ is the ‘normal triad’ of the source with respect to the ICRS (Murray 1983). In this triad, \mathbf{r}_i is the barycentric coordinate direction to the source at time t_{ep} , $\mathbf{p}_i = \langle \mathbf{Z} \times \mathbf{r}_i \rangle$, and $\mathbf{q}_i = \mathbf{r}_i \times \mathbf{p}_i$. The components of these unit vectors in the ICRS are given by the columns of the matrix

$$\mathbf{C}'[\mathbf{p}_i \ \mathbf{q}_i \ \mathbf{r}_i] = \begin{bmatrix} -\sin \alpha_i & -\sin \delta_i \cos \alpha_i & \cos \delta_i \cos \alpha_i \\ \cos \alpha_i & -\sin \delta_i \sin \alpha_i & \cos \delta_i \sin \alpha_i \\ 0 & \cos \delta_i & \sin \delta_i \end{bmatrix}. \quad (5)$$

$\mathbf{b}_G(t)$ is the barycentric position of Gaia at the time of observation, and A_u the astronomical unit. t_B is the barycentric time obtained by correcting the time of observation for the Römer delay; to sufficient accuracy it is given by

$$t_B = t + \mathbf{r}'_i\mathbf{b}_G(t)/c, \quad (6)$$

where c is the speed of light.

In Eq. (4) the term containing μ_{ri} accounts for the relative rate of change of the barycentric distance to the source. This term may produce secular variations of the proper motions and parallaxes of some nearby stars, which in principle allow their radial velocities to be determined astrometrically (Dravins et al. 1999). However, for the vast majority of these stars, μ_{ri} can be more accurately calculated by combining the spectroscopically measured radial velocities with the astrometric parallaxes. Thus, although all six astrometric parameters are taken into account when computing the coordinate direction, usually only five of them are considered as unknowns in the astrometric solution.

The standard model can be derived by considering the uniform motion of the source in a purely classical framework, using




Fig. 2. The Scanning Reference System (SRS) $[x \ y \ z]$ is defined by the viewing directions \mathbf{f}_P and \mathbf{f}_F according to Eq. (2). The instrument angles (φ, ζ) are the spherical coordinates in the SRS of the observed (proper) direction \mathbf{u} towards the object. See Fig. 4 for further definition of the viewing directions. The field sizes are greatly exaggerated in this drawing; in reality the astrometrically useful field is $0.72^\circ \times 0.69^\circ$ (along \times across scan) for each viewing direction. The basic angle is $\Gamma_c = \arccos(\mathbf{f}_P' \mathbf{f}_F)$, nominally equal to 106.5° .

Euclidean coordinates and neglecting the light propagation time from the source to the observer (except for the Römer delay). To the accuracy of Gaia, relativistic and light-propagation effects are by no means negligible, but it can be shown that this model is nevertheless accurate enough to *model* the observations to sub-microarcsec accuracy. It is adopted for this purpose because it closely corresponds to the conventional interpretation of the astrometric parameters. However, when the astrometric parameters are to be *interpreted* in terms of the barycentric space velocity of the source, some of these effects may come into play (Lindegren & Dravins 2003).

The transformation from $\tilde{\mathbf{u}}_i(t)$ to the observable (proper) direction $\mathbf{u}_i(t)$ involves taking into account gravitational light deflection by solar-system bodies and the Lorentz transformation to the co-moving frame of the observer (stellar aberration); the relevant formulae are given by Klioner (2003). This transformation therefore depends also on the global parameters \mathbf{g} , for example the PPN parameter γ which measures the strength of the gravitational light deflection. The calculation uses some auxiliary data, not subject to improvement by the solution, and which are here denoted \mathbf{h} ; normally they include for example the barycentric ephemerides of the Gaia satellite and of solar-system bodies, along with their masses. The complete transformation can therefore be written symbolically as:

$$\mathbf{u}_i(t) = \mathbf{u}(s_i, \mathbf{g} | t, \mathbf{h}), \quad (7)$$

where the vertical bar formally separates the (updatable) parameters s_i and \mathbf{g} from the (fixed) given data t and \mathbf{h} . Strictly speaking, the function \mathbf{u} returns the coordinates of the proper direction in the CoMRS, that is the column matrix $\mathbf{C}'\mathbf{u}_i(t)$. The source parameter vector \mathbf{s} is the concatenation of the sub-vectors s_i for all the primary sources.

3.3. Attitude model

The attitude specifies the instantaneous orientation of the Gaia instrument in the celestial reference frame, that is the transfor-

mation from $\mathbf{C} = [X \ Y \ Z]$ (more precisely, the CoMRS) to $\mathbf{S} = [x \ y \ z]$ (Sect. 3.1) as a function of time. The spacecraft is controlled to follow a specific attitude as a function of time, for example the so-called Nominal Scanning Law (NSL; de Bruijne et al. 2010), designed to provide good coverage of the whole celestial sphere while satisfying a number of other requirements. The NSL is analytically defined for arbitrarily long time intervals by just a few free parameters.

However, the actual attitude will deviate from the intended (nominal) scanning law by up to ~ 1 arcmin in all three axes, and these deviations vary on time scales from seconds to minutes depending on the level of physical perturbations and the characteristics of the real-time attitude measurements and control law (cf. Appendix D.4). The CCD integration time of (usually) 4.42 s means that the true (physical) attitude cannot be observed, but only a smoothed version of it, corresponding to the convolution of the physical attitude with the CCD exposure function (Appendix D.3). This ‘effective’ attitude must however be a posteriori estimable, at any instant, to an accuracy compatible with the astrometric goals, i.e., at sub-mas level. For this purpose the effective attitude is modelled in terms of a finite number of attitude parameters, for which it is necessary to choose a suitable mathematical representations of the instantaneous transformation, as well as of its time dependence. For the Gaia data processing, we have chosen to use quaternions (Appendix A) for the former, and B-splines (Appendix B) for the latter representation.

At any time the orientation of the SRS (\mathbf{S}) with respect to the CoMRS (\mathbf{C}) may be represented by the attitude matrix

$$\mathbf{A} \equiv \begin{bmatrix} A_{xx} & A_{xy} & A_{xz} \\ A_{yx} & A_{yy} & A_{yz} \\ A_{zx} & A_{zy} & A_{zz} \end{bmatrix} = \mathbf{S}'\mathbf{C} = \begin{bmatrix} \mathbf{x}'\mathbf{X} & \mathbf{x}'\mathbf{Y} & \mathbf{x}'\mathbf{Z} \\ \mathbf{y}'\mathbf{X} & \mathbf{y}'\mathbf{Y} & \mathbf{y}'\mathbf{Z} \\ \mathbf{z}'\mathbf{X} & \mathbf{z}'\mathbf{Y} & \mathbf{z}'\mathbf{Z} \end{bmatrix}, \quad (8)$$

which is a proper orthonormal matrix, $\mathbf{A}\mathbf{A}' = \mathbf{I}$, $\det(\mathbf{A}) = +1$. This is also called the direction cosine matrix, since the rows are the direction cosines of x , y and z in the CoMRS, and the columns are the direction cosines of X , Y and Z in the SRS.

The orthonormality of \mathbf{A} implies that the matrix elements satisfy six constraints, leaving three degrees of freedom for the attitude representation. Although one could parametrize each of the nine matrix elements as a continuous function of time, for example using splines, it would in practice not be possible to guarantee that the orthonormality constraints hold at any time. The adopted solution is to represent the instantaneous attitude by a unit quaternion \mathbf{q} , which has only four components, $\{q_x, q_y, q_z, q_w\}$, satisfying the single constraint $q_x^2 + q_y^2 + q_z^2 + q_w^2 = 1$. This is easily enforced by normalization.

The attitude quaternion $\mathbf{q}(t)$ gives the rotation from the CoMRS (\mathbf{C}) to the SRS (\mathbf{S}). Using quaternions, their relation is defined by the transformation of the coordinates of the arbitrary vector \mathbf{v} in the two reference systems,

$$\{\mathbf{S}'\mathbf{v}, 0\} = \mathbf{q}^{-1} \{\mathbf{C}'\mathbf{v}, 0\} \mathbf{q}. \quad (9)$$

In the terminology of Appendix A.3 this is a frame rotation of the vector from \mathbf{C} to \mathbf{S} . The inverse operation is $\{\mathbf{C}'\mathbf{v}, 0\} = \mathbf{q} \{\mathbf{S}'\mathbf{v}, 0\} \mathbf{q}^{-1}$.

Using the B-spline representation in Appendix B, we have

$$\mathbf{q}(t) = \left\langle \sum_{n=\ell-M+1}^{\ell} \mathbf{a}_n B_n(t) \right\rangle, \quad (10)$$

where \mathbf{a}_n ($n = 0 \dots N-1$) are the coefficients of the B-splines $B_n(t)$ of order M defined on the knot sequence $\{\tau_k\}_{k=0}^{N+M-1}$. The function $B_n(t)$ is non-zero only for $\tau_n < t < \tau_{n+M}$, which is

why the sum in Eq. (10) only extends over the M terms ending with $n = \ell$. Here, ℓ is the so-called *left index* of t satisfying $\tau_\ell \leq t < \tau_{\ell+1}$. The normalization operator $\langle \cdot \rangle$ guarantees that $\mathbf{q}(t)$ is a unit quaternion for any t in the interval $[\tau_{M-1}, \tau_N]$ over which the spline is completely defined. Although the coefficients \mathbf{a}_n are formally quaternions, they are not of unit length. The attitude parameter vector \mathbf{a} consists of the components of the whole set of quaternions \mathbf{a}_n .

Cubic splines ($M = 4$) are currently used in this attitude model. Each component of the quaternion (before the normalization in Eq. 10) is then a piecewise cubic polynomial with continuous value, first, and second derivative for any t ; the third derivative is discontinuous at the knots. When initializing the solution, it is possible to select any desired order of the spline. Using a higher order, such as $M = 5$ (quartic) or 6 (quintic), provides improved smoothness but also makes the spline fitting less local, and therefore more prone to undesirable oscillatory behaviour. The flexibility of the spline is in principle only governed by the number of degrees of freedom (that is, in practice, the knot frequency), and is therefore independent of the order. One should therefore not choose a higher order than is warranted by the smoothness of the actual effective attitude, which is difficult to model a priori (cf. Appendix D.4). Determining the optimal order and knot frequency may in the end only be possible through analysis of the real mission data.

Equation (7) returns the observed direction to the source relative to the celestial reference system, or $\mathbf{C}'\mathbf{u} = [u_x, u_y, u_z]'$. In order to compute the position of the image in the field of view, we need to express this direction in the Scanning Reference System, SRS (Sect. 3.1), or $\mathbf{S}'\mathbf{u} = [u_x, u_y, u_z]'$. The required transformation is obtained by a frame rotation according to Eq. (9),

$$\{\mathbf{S}'\mathbf{u}, 0\} = \mathbf{q}^{-1} \{\mathbf{C}'\mathbf{u}, 0\} \mathbf{q}, \quad (11)$$

whereupon the instrument angles (φ, ζ) are obtained from

$$\mathbf{S}'\mathbf{u} \equiv \begin{bmatrix} u_x \\ u_y \\ u_z \end{bmatrix} = \begin{bmatrix} \cos \zeta \cos \varphi \\ \cos \zeta \sin \varphi \\ \sin \zeta \end{bmatrix} \Leftrightarrow \begin{cases} \varphi = \text{atan2}(u_y, u_x) \\ \zeta = \text{atan2}(u_z, \sqrt{u_x^2 + u_y^2}) \end{cases} \quad (12)$$

(Fig. 2), where we adopt the convention that $-\pi \leq \varphi < \pi$. The field index f and the along-scan field angle η are finally obtained as

$$f = \text{sign}(\varphi), \quad \eta = \varphi - f\Gamma_c/2, \quad (13)$$

where the basic angle, Γ_c , is here a purely conventional value (as suggested by the subscript). The field-of-view limitations imply that $|\eta| \leq 0.5^\circ$ and $|\zeta| \leq 0.5^\circ$ for any observation.

Given the time of observation t and the set of source parameters \mathbf{s} , attitude parameters \mathbf{a} , and global parameters \mathbf{g} , the field index f and the field angles (η, ζ) can thus be computed by application of Eqs. (7), (10)–(13). The resulting *computed* field angles are concisely written $\eta(t | \mathbf{s}, \mathbf{a}, \mathbf{g})$, $\zeta(t | \mathbf{s}, \mathbf{a}, \mathbf{g})$.

3.4. Geometric instrument model

The geometric instrument model defines the precise layout of the CCDs in terms of the field angles (η, ζ) . This layout depends on several different factors, including: the physical geometry of each CCD; its position and alignment in the focal-plane assembly; the optical system including its scale value, distortion and other aberrations; and the adopted (conventional) value of the basic angle Γ_c . Some of these (notably the optical distortion) are different in the two fields of view and may evolve slightly over the mission; on the other hand these variations are expected to be

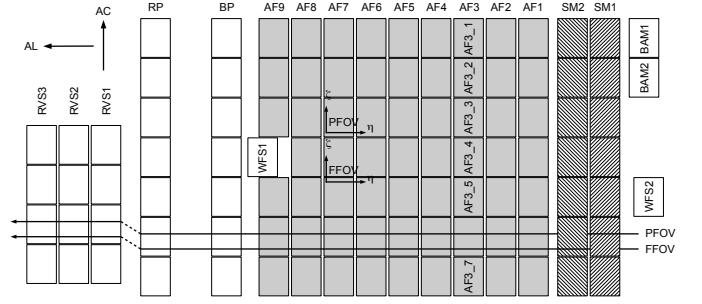


Fig. 3. Layout of the CCDs in Gaia's focal plane. The star images move from right to left in the diagram. The along-scan (AL) and across-scan (AC) directions are indicated in the top left corner. The skymappers (SM1, SM2) provide source image detection, two-dimensional position estimation and field-of-view discrimination. The astrometric field (AF1–AF9) provides accurate AL measurements and (sometimes) AC positions. Additional CCDs used in the blue and red photometers (BP, RP), the radial-velocity spectrometer (RVS), for wavefront sensing (WFS) and basic-angle monitoring (BAM) are not further described in this paper. One of the rows (AF3) illustrates the system for labelling individual CCDs (AF3_1, etc). The nominal paths of two star images, one in the preceding field of view (PFOV) and one in the following field of view (FFOV) are indicated. For purely technical reasons the origin of the field angles (η, ζ) corresponds to different physical locations on the CCDs in the two fields. The physical size of each CCD is $45 \times 59 \text{ mm}^2$.

very smooth as a function of the field angles. Other factors, such as the detailed physical geometry of the pixel columns, may be extremely stable but possibly quite irregular on a small scale. As a result, the geometric calibration model must be able to accommodate both smooth and irregular patterns evolving on different time scales in the two fields of view.

In the following it should be kept in mind that the astrometric instrument of Gaia is optimised for one-dimensional measurements in the along-scan direction, i.e., of the field angle η , while the requirements in the perpendicular direction (ζ) are much relaxed. This feature of Gaia (and Hipparcos) is a consequence of fundamental considerations derived from the requirements to determine a global reference frame and absolute parallaxes (Lindegren & Bastian 2011). In principle the measurement accuracy in the ζ direction, as well as the corresponding instrument modelling and calibration accuracies, may be relaxed by up to a factor given by the inverse angular size of the field of view (Lindegren 2005), or almost two orders of magnitude for a field of 0.7° . In practice a ratio of the order of 10 may be achieved, in which case the across-scan measurement and calibration errors have a very marginal effect on the overall astrometric accuracy.

The focal plane of Gaia contains a total of 106 CCDs (Laborie et al. 2007), of which only the 14 CCDs in the skymappers (SM1 and SM2) and the 62 in the astrometric field (AF1 through AF9) are used for the astrometric solution (Fig. 3). The CCDs have 4500 pixel lines in the along-scan (AL, constant ζ , decreasing η) direction and 1966 pixel columns in the across-scan (AC, constant η , increasing ζ) direction. They are operated in the Time, Delay and Integrate (TDI) mode (also known as drift-scanning), an imaging technique well-known in astronomy from ground-based programmes such as the Sloan Digital Sky Survey (Gunn et al. 1998). Effectively, the charges are clocked along the CCD columns at the same (average) speed as the mo-

tion of the optical images, i.e., 60 arcsec s⁻¹ for Gaia. The exposure (integration) time is thus set by the time it takes the image to move across the CCD, or nominally $T \approx 4.42$ s, if no gate is activated. For this exposure time the central pixels will be saturated for sources brighter than magnitude $G \approx 12$.¹ Gate activation, or ‘gating’ for short, is the adopted method to obtain valid measurements of brighter sources. Gating temporarily inhibits charge transfer across a certain TDI line (row of pixels AC), thus effectively zeroing the charge image and reducing the exposure time in proportion to the number of TDI lines following the gate. A range of discrete exposure times is thus available, the shortest one, according to current planning, using only 16 TDI lines (15.7 ms). Gate selection is made by the on-board software based on the measured brightness of the source in relation to the calibrated full-well capacity of the relevant section of the CCD. Measurement errors and the spread in full-well capacities make the gate-activation thresholds somewhat fuzzy, and a given bright source is not necessarily always observed using the same gate. Moreover, a quasi-random selection of the observations of fainter sources will also be gated, viz., if they happen to be read out at about the same time, and on the same CCD, as a gated bright star. The sky-mappers are operated in a permanently gated mode, so that in practice only the last 2900 TDI lines are used in SM1 and SM2.

The sky-mappers are crucial for the real-time operation of the instrument, since they detect the sources as they enter the field of view, and allow determination of an approximate two-dimensional position of the images and (together with data from AF1) their speed across the CCDs. This information is used by the on-board attitude computer, in order to determine which CCD pixels should be read out and transmitted to ground (Sect. 3.5). The sky-mappers also allow to discriminate between the two viewing directions, since sources in the preceding field of view (PFoV) are only seen by SM1 and sources in the following field of view (FFoV) are only seen by SM2. In the astrometric field (as well as in BP, RP and RVS) the two fields of view are superposed.

Most observations acquired in the astrometric field (AF) are purely one-dimensional through the on-chip AC binning of data (Sect. 3.5). However, sources brighter than magnitude $G \approx 13$ are always observed as two-dimensional images, providing accurate information about the AC pixel coordinate (μ) as well. Some randomly selected fainter images are also observed two-dimensionally (‘Calibration Faint Stars’, CFS). The bright sources, CFS and SM data provide the AC information necessary for the on-ground three-axis attitude determination.

Because of the TDI mode of observation, AL irregularities of the pixel geometry are smeared out and need not be calibrated, and any measurement of the AL or AC position must effectively be referred back to an ‘observation time’ half-way through the integration. Correspondingly, the pixel geometry can be represented by a fiducial ‘observation line’ $[\eta(\mu), \zeta(\mu)]$ traced out in the field angles η, ζ as functions of the AC pixel coordinate μ (Fig. 4). The AC pixel coordinate is defined as a continuous variable with $\mu = 14$ when the image is centred on the first pixel column, $\mu = 15$ at the second pixel column, and so on until $\mu = 1979$ at the last (1966th) pixel column. The offset by 13 pixels allows for the presence of pre-scan pixel data in some observations.

¹ G is the broad-band magnitude measured on Gaia’s unfiltered CCDs. $G \approx V$ for unreddened early-type stars, while $G \approx V - 2$ for M stars. See Jordi et al. (2010) for a precise characterization of G .




Fig. 4. Schematic illustration of how the field angles (η, ζ) are defined in terms of the CCD layout in Fig. 3. For simplicity only the projections of two CCDs, AF8_3 and AF8_4, into the Scanning Reference System (SRS) are shown (not to scale). The field angles have their origins at the respective viewing direction f_P or f_F (Fig. 2), which are defined in relation to the nominal centres of the CCDs (crosses); the actual configuration of the detectors is described by fiducial observation lines according to Eq. (15). The dashed curve shows the apparent path of a stellar image across the two fields of view. Its intersection with the observation lines define the instants of observations.

Nominally, the observation line corresponds to the $(K/2)$ th pixel line projected backwards through the optical instrument onto the Scanning Reference System on the sky, where K is the number of active AL pixel lines in the observation (normally $K = 4500$).² For a gated observation K is much smaller and the observation line is therefore correspondingly displaced towards the CCD readout register. A separate set of geometrical calibration parameters is therefore needed for each gate being used. In the calibration updating (Sect. 5.3) all the calibration parameters are however solved together, with the overlap due to the fuzzy gate-activation thresholds providing the necessary connection between the different gates.

Let n be an index identifying the different CCDs used for astrometry, i.e., for each of the 76 CCDs in the SM and AF part of the focal plane. Furthermore, let g be a gate index such that $g = 0$ is used for non-gated observations and $g = 1, \dots, 12$ for gated observations of progressively brighter sources. In each field of view (index f) the nominal observation lines $[\eta_{fng}^0(\mu), \zeta_{fng}^0(\mu)]$ could in principle be calculated from the nominal characteristics of the focal plane assembly (FPA) and ray tracing through the nominal optical system. However, since the nominal observation lines are only used as a reference for the calibration of the actual observation lines, a very simplistic transformation from linear coordinates to angles can be used without introducing approximation errors in the resulting calibration model. The nominal observation lines are therefore defined by the transformation

$$\left. \begin{aligned} \eta_{fng}^0(\mu) &\equiv \eta_{ng}^0 = -Y_{FPA}[n, g]/F \\ \zeta_{fng}^0(\mu) &= -(X_{FPA}[n] - (\mu - \mu_c)p_{AC} - X_{FPA}^{\text{centre}}[f])/F \end{aligned} \right\} \quad (14)$$

where X_{FPA}, Y_{FPA} are physical coordinates (in mm) in the focal plane,³ $X_{FPA}[n]$ is the physical AC coordinate of the nom-

² In reality the definition of the fiducial observation line is a bit more complex, as some of the pixel lines are blocked out by an aluminium mask.

³ For consistency with notations adopted by the ESA project team and the industrial contractor, and extensively used, e.g., for on-ground calibrations, $+X_{FPA}$ is oriented along $-\zeta$ and $+Y_{FPA}$ along $-\eta$.

inal centre of the n th CCD, $Y_{\text{FPA}}[n, g]$ the physical AL coordinate of the nominal observation line for gate g on the n th CCD, and $X_{\text{FPA}}^{\text{centre}}[f]$ the physical AC coordinate of the nominal field centre for field index f . $\mu_c = 996.5$ is the AC pixel coordinate of the CCD centre, $p_{\text{AC}} = 30 \mu\text{m}$ the physical AC pixel size, and $F = 35$ m the nominal equivalent focal length of the telescope. While η_{ng}^0 is independent of f , in the AC direction the origins are offset by about ± 221 arcsec between the two fields of view, corresponding to the physical coordinates $X_{\text{FPA}}^{\text{centre}}[f] = (-37.5 \text{ mm})f$.

It is emphasized that the nominal observation lines are purely conventional reference quantities, and need not be recomputed, e.g., once a more accurate estimate of F becomes available.

Because of possible changes in the instrument during the mission, the actual observation lines will be functions of time. The time dependence is quantified by introducing independent sets of calibration parameters for successive, non-overlapping time intervals. Different groups of parameters may develop on different time scales, and the resulting formulation can be quite complex. For the sake of illustration, let us distinguish between three categories of geometric calibration parameters:

1. Large-scale AL calibration, $\Delta\eta$. This may (minutely but importantly) change due to thermal variations in the optics, the detectors, and their supporting mechanical structure. These variations could occur on short time scales (of the order of a day), and would in general be different in the two fields of view. They are modelled as low-order polynomials in the across-scan pixel coordinate μ .
2. Small-scale AL calibration, $\delta\eta$. This is mainly related to physical defects or irregularities in the CCDs themselves, for example ‘stitching errors’ introduced by the photolithography process used to manufacture the CCDs. These irregularities are expected to be stable over very long time scales, possibly throughout the mission, and should be identical in both fields of view. They require a spatially detailed modelling, perhaps with a resolution of just one or a few AC pixels.
3. Large-scale AC calibration, $\Delta\zeta$. Although the physical origin is the same as for $\Delta\eta$, the AC components can be assumed to be constant on long time scales, since the calibration requirement in the AC direction is much more relaxed than in the AL direction. They are modelled as low-order polynomials in the field angles, separately in each field of view.

Let index j identify the ‘short’ time intervals needed for the large-scale AL calibration, and index k identify the ‘long’ time intervals applicable to the small-scale AL calibration and large-scale AC calibration. That is, an observation at time t belongs to some short time interval j and some long time interval k , where j and k are readily computed from the known t and the corresponding sequences of division times.⁴ Assuming that quadratic polynomials in μ are sufficient for the large-scale calibrations, and that full AC pixel resolution is required for the small-scale AL calibration, the observation lines at time t are modelled as

$$\left. \begin{aligned} \eta_{fng}(\mu, t) &= \eta_{ng}^0 + \sum_{r=0}^2 \Delta\eta_{rfngj} L_r^* \left(\frac{\mu - 13.5}{1966} \right) + \delta\eta_{ngkm} \\ \zeta_{fng}(\mu, t) &= \zeta_{fng}^0(\mu) + \sum_{r=0}^2 \Delta\zeta_{rfngk} L_r^* \left(\frac{\mu - 13.5}{1966} \right) \end{aligned} \right\} \quad (15)$$

⁴ Technically, the use of independent parameter values in successive time intervals represents a spline of order 1 (i.e., degree 0), the separation times constitute the knot sequence, and j or k correspond to the ‘left index’ (Appendix B.2).

where $L_r^*(x)$ is the shifted Legendre polynomial of degree r (orthogonal on $[0, 1]$), i.e., $L_0^*(x) = 1$, $L_1^*(x) = 2x - 1$, $L_2^*(x) = 6x^2 - 6x + 1$, etc.; $\Delta\eta_{rfngj}$ are the large-scale AL parameters, $\delta\eta_{ngkm}$ the small-scale AL parameters (with $m = \text{round}(\mu)$ the index of the nearest pixel column), and $\Delta\zeta_{rfngk}$ the large-scale AC parameters.

In order to render all the geometric calibration parameters uniquely determinable, a number of constraints are rigorously enforced by the astrometric solution. Effectively, they define the origins of the field angles, i.e., the viewing directions \mathbf{f}_P and \mathbf{f}_F , to coincide with the average nominal field angles of the CCD centres. The necessary constraints are:

$$\sum_f \sum_{n \in \text{AF}} \Delta\eta_{0fn0j} = 0 \quad \text{for each } j, \quad (16)$$

$$\sum_m \delta\eta_{n0km} L_r^* \left(\frac{m - 13.5}{1966} \right) = 0 \quad \text{for each combination } rnk, \quad (17)$$

$$\sum_{n \in \text{AF}} \Delta\zeta_{0fn0k} = 0 \quad \text{for each combination } fk, \quad (18)$$

Note that $g = 0$ throughout in Eq. (16)–(18). That the constraints are only defined in terms of the non-gated observations ($g = 0$) implies that the observation lines for the gated observations must be calibrated relative to the non-gated observations. This is possible since any given bright source will not always be observed with the same gate.

Constraint (16) effectively defines the zero point of the AL field angle η by requiring that $\overline{\Delta\eta} = 0$ when averaging over the CCDs and between the two fields of view. Constraint (17) implies that the small-scale AL corrections $\delta\eta$ do not have any components that could be described by $\Delta\eta$ instead; it therefore ensures a unique division between the large-scale and small-scale components. Constraint (18) effectively defines the zero point of the AC field angle ζ by requiring that $\overline{\Delta\zeta} = 0$, separately in each field of view, when averaged over the CCDs. The sums over n in Eqs. (16) and (18) are restricted to the CCDs in the astrometric field (AF), since the skymapper (SM) measurements are less accurate.

The basic angle Γ_c introduced in Sect. 3.3 is a fixed reference value, and any real variations of the angle between the two lines of sight will therefore show up as a variation in $\Delta\eta$ with opposite signs in the two fields of view. The offset of the actual basic angle with respect to the conventional value Γ_c may be defined as the average difference of $\Delta\eta$ between the two fields of view, where the average is computed over the astrometric CCDs, for gate $g = 0$, and over μ ; the result for time period j is

$$\Delta\Gamma_j = \frac{1}{62} \sum_{n \in \text{AF}} \sum_f f \Delta\eta_{0fn0j}. \quad (19)$$

Although this is obviously a useful quantity to monitor, it does not appear as a parameter in the geometric calibration model. A number of additional quantities representing the mean scale offset, the mean field orientation, etc., can similarly be computed from the large-scale calibration parameters for the purpose of monitoring.

Equation (15) encapsulates a specific formulation of the geometric instrument model, with certain assumptions about the shape, dependencies, and time scales of possible variations. While this particular model is currently believed to be sufficient to describe the behaviour of the actual instrument to the required accuracy, it is very likely that modifications will be needed after a first analysis of the flight data. Moreover, in the course of

the data analysis one may want to try out alternative models, or examine possible systematics resulting from the pre-processing (location estimator). In order to facilitate this, a much more flexible *generic calibration model* has been implemented. In the generic model, the ‘observed’ field angles (representing the true observation lines) for any observation l are written

$$\left. \begin{aligned} \eta_l^{\text{obs}} &= \eta_{ng}^0 + \sum_{r=0}^{N_{\text{AL}}-1} E_r^{\text{AL}}(l) \\ \zeta_l^{\text{obs}} &= \zeta_{fng}^0(\mu) + \sum_{r=0}^{N_{\text{AC}}-1} E_r^{\text{AC}}(l) \end{aligned} \right\} \quad (20)$$

where each of the $E_r(l)$ (for brevity dropping superscript AL/AC) represents a basic calibration effect, being a linear combination of elemental calibration functions Φ_{rs} :

$$E_r(l) = \sum_{s=0}^{K_r-1} c_{rs} \Phi_{rs}(l). \quad (21)$$

N_{AL} and N_{AC} are the number of effects considered along and across scan. The whole set of coefficients c_{rs} forms the calibration vector \mathbf{c} .

In the generic formulation, the multiple indices f , n , g and the variables μ and t are replaced by the single observation index l . This allows maximum flexibility in terms of how the calibration model is implemented in the software. The functions Φ_{rs} receive the observation index l , and it is assumed that this index suffices to derive from it all quantities needed to evaluate the calibration functions for this observation. Examples of quantities that can be derived from the observation index are: the FoV index f , the CCD and gate indices n and g , the AC pixel coordinate μ , time t , and any relevant astrometric, photometric or spectroscopic parameter of the source (magnitude, colour index, etc.). Intrinsically real-valued quantities such as t can be subdivided into non-overlapping intervals with different sets of calibration parameters applicable to each interval. The basic calibration model (15) can therefore be implemented as a particular instance of the generic model, with for example $L_r^*(x)$ ($r = 0, 1, 2$) constituting three of the calibration functions (with μ derived from l). Once the calibration functions have been coded, the entire calibration model can be conveniently specified (and changed) through an external configuration file alone using, e.g., XML structures.

Some elemental calibration functions may be introduced for diagnostic purposes rather than actual calibration. An example of this is any function depending on the colour or magnitude of the source. The origin of such effects is briefly explained in Appendix D.1 and D.2. Magnitude- and colour-dependent variations of the instrument response are expected to be fully taken into account by the signal modelling on the CCD data level, as outlined in Sect. 3.5 (notably by the LSF and CDM calibrations), and should not show up in the astrometric solution. Thus, non-zero results for such ‘non-geometric’ diagnostic calibration parameters indicate that the signal modelling should be improved. In the final solution the diagnostic calibration parameters should ideally be zero.

The parameters of the generic calibration model must satisfy a number of constraints similar to Eqs. (16)–(18). These can be cast in the general form

$$\mathbf{C}'\mathbf{c} = \mathbf{0}, \quad (22)$$

where the matrix \mathbf{C} contains one column, with known coefficients, for each constraint. The columns are, by design, linearly independent.

3.5. Signal (CCD data) model

As suggested in Fig. 1, the modelling of CCD data at the level of individual pixels (i.e., the photon counts) is not part of the geometrical model of the observations with which we are concerned in this paper. However, the processing of the photon counts, effectively by fitting the CCD data model, produces the ‘observations’ that are the input to the astrometric core solution. In order to clarify the exact meaning of these observations we include here a brief overview of the signal model.

The pixel size, $10 \mu\text{m} \approx 59 \text{ mas}$ in the along-scan (AL) direction and $30 \mu\text{m} \approx 177 \text{ mas}$ in the across-scan (AC) direction, roughly matches the theoretical diffraction image for the $1.45 \times 0.50 \text{ m}^2$ telescope pupil of Gaia (effective wavelength $\sim 650 \text{ nm}$). Around each detected object, only a small rectangular window (typically 6–18 pixels long in the AL direction and 12 pixels wide in the AC direction) is actually read out and transmitted to the ground. Moreover, for most of the observations in the astrometric field (AF), on-chip binning in the serial register is used to sum the charges over the window in the AC direction. This effectively results in a one-dimensional image of 6–18 AL ‘samples’, where the signal (N_k) in each sample k is the sum of 12 AC pixels. The exact time t_k when a sample was transferred to the serial register, expressed on the TCB scale, is in principle known from the time correlation of the on-board clock with ground-based time signals. Because of the known one-to-one relation between the TDI period counter k and t_k , we may use k as a proxy for t_k in subsequent calculations.

For single stars, and in the absence of the effects discussed in Appendix D.2, the sample values in the window are modelled as a stochastic variable (Poisson photon noise plus electronic read-out noise) with expected value

$$\lambda_k \equiv \mathbb{E}(N_k) = \beta + \alpha L(k - \kappa) \quad (23)$$

where β , α and κ are the so-called image parameters representing the (uniform) background level, the amplitude (or flux) of the source, and the AL location (pixel coordinate) of the image centroid. The continuous, non-negative function $L(x)$ is the Line Spread Function (LSF) appropriate for the observation. $L(x)$ depends, for example, on the spectrum of the source and on the position of the image in the focal plane. The (in general non-integer) pixel coordinate κ is expressed on the same scale as the (integer) TDI index k , and may be translated to the equivalent TCB $t(\kappa)$ by means of the known relation between k and t_k . The *CCD observation time* is defined as $t(\kappa - K/2)$, where K is the number of TDI periods used for integrating the image (see Appendix D.3). Formally, the CCD observation time is the instant of time at which the centre of the source image passed across the CCD fiducial line halfway between the first and the last TDI line used in the integration (this will depend on the gating).

Fitting the CCD signal model to the one-dimensional sample values N_k thus gives, as the end result of observation l , an estimate of the AL pixel coordinate κ of the image in the pixel stream, which is then transformed to the observation time t_l . The fitting procedure also provides an estimate of the uncertainty in κ , which can be expressed in angular measure as a formal standard deviation of the AL measurement, σ_l^{AL} . It is derived by error propagation through the fitted signal model, taking into account the dominant noise sources, photon noise and readout noise.

For some observations, AC information is also provided through two-dimensional sampling of the pixel window around

the detected object. This applies to all SM observations, AF observations of bright ($G \lesssim 13$) sources, and AF observations of Calibration Faint Stars (Sect. 3.4). The modelling of the two-dimensional images follows the same principles as outlined above for the one-dimensional (AL only) case, only that the LSF is replaced by a two-dimensional Point Spread Function (PSF) and that there is one more location parameter to estimate, namely the AC pixel coordinate μ . The astrometric result in this case consists of the observation time t_l , the observed AC coordinate μ_l , and their formal standard uncertainties σ_l^{AL} and σ_l^{AC} (both expressed as angles).

The estimation errors for different images are, for the subsequent analysis, assumed to be statistically independent (and therefore uncorrelated). This is a very good approximation to the extent that they only depend on the photon and readout noises. However, modelling errors at the various stages of the processing (in particular CTI effects in the signal modelling [Appendix D.3] and attitude irregularities [Appendix D.4]) are likely to introduce errors that are correlated at least over the nine AF observations in a field transit. The resulting correlations as such are not taken into account in the astrometric solution (i.e., the weight matrix of the least-squares equations is taken to be diagonal), but the sizes of the modelling errors are estimated, and are employed to reduce the statistical weights of the observations as described in Sect. 3.6. The AL and AC estimates of a given (two-dimensional) image are roughly independent at least in the limit of small optical aberrations.

3.6. Synthesis model

By synthesis of the models described in the preceding sections, we are now in a position to formulate very precisely the core astrometric data analysis problem as outlined in Sect. 2. The unknowns are represented by the vectors \mathbf{s} , \mathbf{a} , \mathbf{c} , and \mathbf{g} of respectively the source, attitude, calibration, and global parameters. For any observation l the observed quantities are the observation time t_l and, where applicable, the observed AC pixel coordinate μ_l , with their formal uncertainties σ_l^{AL} , σ_l^{AC} . We then have the global minimization problem

$$\min_{\mathbf{s}, \mathbf{a}, \mathbf{c}, \mathbf{g}} Q = \sum_{l \in \text{AL}} \frac{(R_l^{\text{AL}})^2 w_l^{\text{AL}}}{(\sigma_l^{\text{AL}})^2 + (\epsilon_l^{\text{AL}})^2} + \sum_{l \in \text{AC}} \frac{(R_l^{\text{AC}})^2 w_l^{\text{AC}}}{(\sigma_l^{\text{AC}})^2 + (\epsilon_l^{\text{AC}})^2}, \quad (24)$$

where

$$R_l^{\text{AL}}(\mathbf{s}, \mathbf{a}, \mathbf{c}, \mathbf{g}) = \eta_{fng}(\mu_l, t_l | \mathbf{c}) - \eta(t_l | \mathbf{s}, \mathbf{a}, \mathbf{g}), \quad (25)$$

$$R_l^{\text{AC}}(\mathbf{s}, \mathbf{a}, \mathbf{c}, \mathbf{g}) = \zeta_{fng}(\mu_l, t_l | \mathbf{c}) - \zeta(t_l | \mathbf{s}, \mathbf{a}, \mathbf{g}) \quad (26)$$

are the residuals in the field angles, taken as functions of the unknowns⁵, and $l \in \text{AL}$ refers to observations with a valid AL component, etc. The applicable indices f , n , g are of course known for every observation l . In Eq. (24), ϵ_l^{AL} and ϵ_l^{AC} represent all AL and AC error sources extraneous to the formal uncertainties; they include in particular modelling errors in the source behaviour (e.g., for unrecognized binaries), attitude and calibration, which have to be estimated as functions of time and source in the course of the data analysis process. w_l^{AL} and w_l^{AC} are weight factors, always in the range 0 to 1; for most observations they are equal to 1, but ‘bad’ data (outliers) are assigned smaller weight factors. The determination of these factors is discussed in Sect. 5.1.2.

⁵ Note that in Eq. (25) the quantity μ_l is just a given value (observed or computed); in the case of one-dimensional images the observed μ_l is replaced by an approximate value derived from current knowledge on the source and attitude.

For the sake of conciseness we shall hereafter consider the AL and AC components of an observation to have separate observation indices l , so that for example R_l stands for either R_l^{AL} or R_l^{AC} , as the case may be. This allows the two sums in Eq. (24) to be contracted and written concisely as

$$Q(\mathbf{s}, \mathbf{a}, \mathbf{c}, \mathbf{g}) = \sum_l \frac{R_l^2 w_l}{\sigma_l^2 + \epsilon_l^2} = \sum_l R_l^2 W_l, \quad (27)$$

where $W_l = w_l/(\sigma_l^2 + \epsilon_l^2)$ is the statistical weight of the observation.

The excess noise ϵ_l represents modelling errors and should ideally be zero. However, it is unavoidable that some sources do not behave exactly according to the adopted astrometric model (Sect. 3.2), or that the attitude sometimes cannot be represented by the spline model in Sect. 3.3 to sufficient accuracy. The excess noise term in Eq. (27) is introduced to allow these cases to be handled in a reasonable way, i.e., by effectively reducing the statistical weight of the observations affected. It should be noted that these modelling errors are assumed to affect *all* the observations of a particular star, or *all* the observations in a given time interval. (By contrast, the downweighting factor w_l is intended to take care of isolated outliers, for example due to a cosmic-ray hit in one of the CCD samples.) This is reflected in the way the excess noise is modelled as the sum of two components,

$$\epsilon_l^2 = \epsilon_i^2 + \epsilon_a^2(t_l), \quad (28)$$

where ϵ_i is the excess noise associated with source i (if $l \in i$, that is, l is an observation of source i), and $\epsilon_a(t)$ is the excess attitude noise, being a function of time. For a ‘good’ primary source, we should have $\epsilon_i = 0$, and for a ‘good’ stretch of attitude data we may have $\epsilon_a(t) = 0$. Calibration modelling errors are not explicitly introduced in Eq. (28), but could be regarded as a more or less constant part of the excess attitude noise. The estimation of ϵ_i is described in Sect. 5.1.2, and the estimation of $\epsilon_a(t)$ in Sect. 5.2.5.

Three separate functions are needed to describe the excess attitude noise, corresponding to AL observations, AC observations in the preceding field of view (ACP), and AC observations in the following field of view (ACF). We distinguish between these functions by letting the subscript a in $\epsilon_a(t)$ stand for either AL, ACP or ACF.

4. Solving the global minimization problem

Assuming 10^8 primary sources, the number of unknowns in the global minimization problem, Eq. (24), is about 5×10^8 for the sources (\mathbf{s}), 4×10^7 for the attitude (\mathbf{a} , assuming a knot interval of 15 s for the 5 yr mission; cf. Sect. 5.2.1), 10^6 for the calibration \mathbf{c} , and less than 100 global parameters (\mathbf{g}). The number of elementary observations (l) considered is about 8×10^{10} . However, the size of the data set, and the large number of parameters, would not by themselves be a problem if the observations, or sources, could be processed sequentially. The difficulty is caused by the strong connectivity among the observations: each source is effectively observed relative to a large number of other sources simultaneously in the field of view, or in the complementary field of view some 106.5° away on the sky, linked together by the attitude and calibration models. The complexity of the astrometric solution in terms of the connectivity between the sources provided by the attitude modelling was analysed by Bombrun et al. (2010), who concluded that a direct solution is infeasible, by

many orders of magnitude, with today's computational capabilities. The study neglected the additional connectivity due to the calibration model, which makes the problem even more unrealistic to attack by a direct method. Note that this impossibility is not a defect, but a virtue of the mathematical system under consideration: it guarantees that a unique, coherent and completely independent global solution for the whole sky can be derived from the system.

The natural alternative to a direct solution is to use some iterative method. This is in fact the standard way to deal with large, sparse systems of equations. The literature in the field is vast and a plethora of methods exist for various kinds of applications. However, the iterative method adopted for Gaia did not spring from a box of ready-made algorithms. Rather, it was developed over several years in parallel with the software system in which it could be implemented and tested. Originally based on an intuitive and somewhat simplistic approach, the algorithm has developed through a series of experiments, insights and improvements into a rigorous, efficient and well-understood procedure, completely adapted to its particular application. In this section we first describe the approach in broad outline, then provide the mathematical background for its better understanding and further development.

4.1. Outline of the iterative solution

The numerical approach to the Gaia astrometry is a block-iterative least-squares solution, referred to as the Astrometric Global Iterative Solution (AGIS). In its simplest form, four blocks are evaluated in a cyclic sequence until convergence. The blocks map to the four different kinds of unknowns outlined in Sect. 3, namely:

- S: the source (star) update, in which the astrometric parameters s of the primary sources are improved;
- A: the attitude update, in which the attitude parameters a are improved;
- C: the calibration update, in which the calibration parameters c are improved;
- G: the global update, in which the global parameters g are improved.

The G block is optional, and will perhaps only be used in some of the final solutions, since the global parameters can normally be assumed to be known a priori to high accuracy.

The blocks must be iterated because each one of them needs data from the three other processes. For example, when computing the astrometric parameters in the S block, the attitude, calibration and global parameters are taken from the previous iteration. The resulting (updated) astrometric parameters are used the next time the A block is run, and so on. This iterative approach to the astrometric solution was proposed early on in the Hipparcos project as an alternative to the 'three-step method' subsequently adopted for the original Hipparcos reductions; see ESA (1997, Vol. 3, p. 488) and references therein. Indeed, the later re-reduction of the Hipparcos raw data by van Leeuwen (2007) used a very similar iterative method, and yielded significantly improved results mainly by virtue of the much more elaborate attitude modelling implemented as part of the approach.

While the block-iterative solution as outlined above is intuitive and appealing in its simplicity, it is from a mathematical standpoint not obvious that it must converge; and if it does indeed converge, it is not obvious how many iterations are required, whether the order of the blocks in each iteration matter,

and whether the converged results do in fact constitute a solution to the global minimization problem. These are fundamental questions for the validity of the adopted iterative approach, and we therefore take some care in the following subsections to establish its theoretical foundations (Sects. 4.2–4.5). Section 5 then describes each of the S, A, C and G blocks in some detail. In addition to these blocks, separate processes are required for the alignment of the astrometric solution with the ICRS, the selection of primary sources, and the calculation of standard uncertainties; these auxiliary processes are discussed in Sect. 6.

4.2. Least-squares approach

Strictly speaking, Eq. (24) is not a least-squares problem, because of the weight factors w_l^{AL} , w_l^{AC} (as well as the excess noises ϵ_l^{AL} , ϵ_l^{AC}), which depend on the AL and AC residuals and hence on the unknowns (s, a, c, g). In Eq. (1) this dependence is formally included in the unspecified metric \mathcal{M} , which therefore is not simply a (weighted) Euclidean norm.

In principle, the minimization problem (24) can be solved by finding a point where the partial derivatives of the objective function Q with respect to all the unknowns are simultaneously zero. In practice, however, the partial derivatives are not computed completely rigorously, and the problem solved is therefore a slightly different one from what is outlined above. In order to understand precisely the approximations involved, it is necessary to consider how different kinds of non-linearities enter the problem.

The functions $\eta_{fng}(\mu, t|c)$ and $\zeta_{fng}(\mu, t|c)$ appearing in Eqs. (25)–(26) are strictly linear in the calibration parameters c , by virtue of the basic geometric calibration model in Eq. (15) or the generic model in Eqs. (20)–(21). On the other hand, the functions $\eta(t|s, a, g)$ and $\zeta(t|s, a, g)$ are non-linear in s, a , and g due to the complex transformations involved (Sects. 3.2–3.3). However, thanks to the data processing prior to the astrometric solution, the initial errors in these parameters are already so small that the corresponding errors in η and ζ are only some 0.1 arcsec ($\sim 10^{-6}$ rad). Second-order terms are therefore typically less than 10^{-12} rad $\simeq 0.2 \mu\text{as}$, that is negligible in comparison with the noise of a single AL observation (some $100 \mu\text{as}$). This means that the partial derivatives of the residuals R_l^{AL} and R_l^{AC} with respect to all the unknowns do not change in the course of the solution. (In practice they are in fact recomputed in each iteration, although that is mainly done because it is more convenient than to store and retrieve the values; nevertheless, this takes care of any remaining non-linearity, however small.) The non-linearities of the underlying astrometric, attitude, calibration and global models are therefore not an issue for the minimization problem as such.

The weight factors w_l represent a different kind of non-linearity, potentially much more important for the final solution. These factors are introduced to make the solution robust against outliers, by reducing their influence on Q and hence on the estimated parameter values (Sect. 5.1.2). Ideally, outlying observations should not contribute at all to the solution (by having $w_l = 0$), while 'normal' observations should receive full weight ($w_l = 1$). In reality there will however be a transition region where the weight factors are between 0 and 1. Since the weight factors are in practice determined by the normalized residuals, $\hat{R}_l \equiv R_l(\sigma_l^2 + \epsilon_l^2)^{-1/2}$, which in turn depend on the parameter values (s, a, c, g), it follows that the partial derivatives contain extra terms of the form $\hat{R}_l^2 \partial w_l / \partial s$, $\hat{R}_l^2 \partial w_l / \partial a$, etc., that are non-zero for some observations. Analogous considerations apply to

the excess noise terms ϵ_l : they too are estimated by means of the residuals (Sects. 5.1.2 and 5.2.5), and therefore in principle introduce additional terms in the partial derivatives.

We take the somewhat pragmatic approach of neglecting the terms depending on the partial derivatives of w_l and ϵ_l with respect to the unknowns when seeking the solution to the global minimization problem. The consequences of this approximation can be appreciated by observing that the down-weighting ($w_l < 1$) only kicks in when the absolute value of the residual exceeds a few times the total standard uncertainty, or some 0.5–1 mas for typical observations of bright sources. Similarly the estimated ϵ_l are only sensitive to changes of the residuals of a similar size. Experience with AGIS runs on simulated data show that the typical changes of the residuals fall below this level already in the first few iterations. Thereafter the weight factors and the estimated excess noises do not change significantly. Neglecting the derivatives of w_l and ϵ_l in the global minimization problem is therefore equivalent to solving the weighted least-squares problem with w_l and ϵ_l fixed at whatever values they settle to after the initial iterations. This is a reasonable assumption given that the statistical weight of any observation, and the size of the modelling errors, are not a priori expected to depend on the actual values of the parameters. The precise solution does of course depend on how w_l and ϵ_l are estimated, but that is an unavoidable consequence of any practical data analysis approach.

4.3. Normal equations

The minimization of Q in Eq. (27) is thus solved by the weighted least-squares method, assuming fixed weights W_l that are however determined as part of the solution. The normal equations for the sources are given by $\frac{1}{2}\partial Q/\partial s = \mathbf{0}$, and similarly for the other unknowns. Linearising around any reference point ($s^{\text{ref}}, a^{\text{ref}}, c^{\text{ref}}, g^{\text{ref}}$) within the linear regime of parameter space, i.e., setting $s = s^{\text{ref}} + x_s$, and similarly for the other unknowns, and expanding to first order in the differential vector $x = (x'_s, x'_a, x'_c, x'_g)'$, we find the normal equations as

$$\left[\sum_l \frac{\partial R_l}{\partial x} \frac{\partial R_l}{\partial x'} W_l \right] x = - \sum_l \frac{\partial R_l}{\partial x} R_l(s^{\text{ref}}, a^{\text{ref}}, c^{\text{ref}}, g^{\text{ref}}) W_l. \quad (29)$$

This can be written in matrix form as

$$N x = b, \quad (30)$$

where N is a symmetric matrix. We now proceed to analyse the structure of this linear system of equations in terms of the previously mentioned block updates S, A, C, G.

The matrix N and the vectors (column matrices) x , b can be partitioned into sub-matrices and sub-vectors corresponding to the different parameter vectors s , a , c , and g :

$$\begin{bmatrix} N_{ss} & N_{sa} & N_{sc} & N_{sg} \\ N_{as} & N_{aa} & N_{ac} & N_{ag} \\ N_{cs} & N_{ca} & N_{cc} & N_{cg} \\ N_{gs} & N_{ga} & N_{gc} & N_{gg} \end{bmatrix} \begin{bmatrix} x_s \\ x_a \\ x_c \\ x_g \end{bmatrix} = \begin{bmatrix} b_s \\ b_a \\ b_c \\ b_g \end{bmatrix}, \quad (31)$$

where $N_{as} = N'_{sa}$, etc. Of importance here is that N_{ss} and N_{aa} have a particularly simple structure. Since s is sub-divided into vectors s_i of length 5 for the individual primary sources (i), it is natural to sub-divide N_{ss} into blocks of 5×5 elements. From Eq. (29) it follows that the (i, j) th such block is given by

$$[N_{ss}]_{ij} = \sum_l \frac{\partial R_l}{\partial s_i} \frac{\partial R_l}{\partial s_j'} W_l = \begin{cases} \sum_{l \in i} \frac{\partial R_l}{\partial s_i} \frac{\partial R_l}{\partial s_i'} W_l & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases} \quad (32)$$

where $l \in i$ signifies that the sum is taken over the observations of source i . The result for $i \neq j$ follows because no observation l is associated with more than one primary source. N_{ss} is consequently block-diagonal, and $N_{ss}^{-1} b_s$ can trivially be computed for arbitrary vector b_s by looping through the sources and solving the corresponding 5×5 system for each source.⁶ This is exactly what is done in the source update block (S).

The vector of attitude unknowns is naturally sub-divided into vectors a_n of length 4, containing the elements of the quaternions \mathbf{a}_n that serve as coefficients in the B-spline representation, Eq. (10). If N_{aa} is correspondingly sub-divided into blocks of 4×4 elements, it follows from Eq. (29) that the (n, m) th such block is given by

$$[N_{aa}]_{nm} = \sum_l \frac{\partial R_l}{\partial a_n} \frac{\partial R_l}{\partial a_m'} W_l. \quad (33)$$

With ℓ denoting the left index of t_l (Sect. 3.3), we have $\partial R_l/\partial a_n = 0$ whenever $n < \ell - M + 1$ or $n > \ell$, where M is the order of the spline. It follows that $[N_{aa}]_{nm} = 0$ if $|n - m| > M - 1$ (cf. Appendix B.1). The non-zero blocks in N_{aa} are therefore confined to the diagonal and $M - 1$ blocks above and below the diagonal (Fig. 5). Thus, $N_{aa}^{-1} b_a$ can be efficiently computed for arbitrary b_a since the Cholesky decomposition of the matrix does not produce any additional fill-in (Appendix C). This system is solved in the attitude update block (A).

In the geometric instrument model (Sect. 3.4), each CCD/gate and time interval combination (index ngk for example in Eq. 15) has its own set of unknowns. Moreover, a given observation l can only refer to one CCD/gate combination. By the same reasoning as above, the sub-matrix N_{cc} is therefore block-diagonal, and $N_{cc}^{-1} b_c$ can be computed for arbitrary b_c by looping over the CCD/gates combinations. This is done in the calibration update block (C). Although the number of calibration parameters per CCD/gate combination can be fairly large ($\sim 10^4$), the resulting systems are well within the bounds that can readily be handled by direct matrix methods, even without taking into account their sparseness.

By definition all the global parameters affect each observation, and the sub-matrix N_{gg} is therefore full. However, since the number of global parameters is never large, $N_{gg}^{-1} b_g$ can easily be computed, which is done in the global update block (G).

In the description above we have implicitly assumed that each of the diagonal blocks N_{ss} , N_{aa} , N_{cc} , and N_{gg} is non-singular, and even well-conditioned in order to avoid numerical instability. This is equivalent to the statement that each of the blocks S, A, C and G is a well-posed problem: for example, that the determination of the source parameters is ‘easy’ if we assume that we know the attitude, calibration and global parameters. This will in practice be guaranteed by the choice of primary sources (which will make $[N_{ss}]_{ii}$ well-conditioned for every i) for the S block, and by the adopted attitude, calibration and global parametrizations, including the constraints necessary to render the updates unique – in particular the quaternion length normalization in Eq. (10) for the attitude model, and Eq. (22) for the calibration model.

Turning now to the off-diagonal sub-matrices of N , it is natural to sub-divide for example the sub-matrix N_{as} into blocks of 4×5 elements corresponding to the 4 components of the quater-

⁶ Here, and elsewhere in this paper, an expression like $A^{-1}b$ is shorthand notation for solving the system $Ay = b$. The inverse matrix A^{-1} is (usually) not computed, but only the solution vector $y = A^{-1}b$ itself.

nion and the 5 astrometric parameters. The (n, i) th block,

$$[N_{as}]_{ni} = \sum_l \frac{\partial R_l}{\partial \mathbf{a}_n} \frac{\partial R_l}{\partial \mathbf{s}'_i} W_l, \quad (34)$$

is non-zero only if source i was observed in the time interval $[\tau_n, \tau_{n+M}]$, which is the support of the B-spline $B_n(t)$. N_{as} is therefore very sparse, but it also has no simple structure because the distribution of the non-zero blocks is linked to the scanning law. Fortunately, as will be explained in Sect. 4.5, there is no need to explicitly compute, much less store, this sub-matrix, nor any of the other off-diagonal sub-matrices in Eq. (31).

4.4. Rank of the normal equations

From the nature of the astrometric observations, which are in effect differential within the (combined) field of view, and the modelling of the primary sources, which does not assume that any of their positions or proper motions are known a priori, it is clear that there is no unique astrometric solution to the problem as outlined above. The fundamental reason for this is that any (small) change in the orientation of the celestial reference system, as well as the introduction of a (small) inertial spin of the system, would leave all observations invariant. In principle the non-uniqueness of the solution is not a problem as such, since the resulting system of positions and proper motions are afterwards aligned with the ICRS by a special process (Sect. 6.1). However, it does imply that the normal matrix N is in principle singular,⁷ which may have consequences for the numerical solution of the normal equations. We say singular ‘in principle’ because arithmetic rounding errors will in practice prevent it from becoming truly singular, although it remains extremely ill-conditioned.

More precisely, we expect N to have rank $n-6$, if $n = \dim(N)$ is the total number of unknowns. The null space of the matrix,

$$N(N) = \{\mathbf{v} \in \mathbb{R}^n \mid N\mathbf{v} = \mathbf{0}\} \quad (35)$$

therefore has rank six. Indeed, it is easy to find six linearly independent vectors $\mathbf{v}_0, \dots, \mathbf{v}_5$ that span the null space: the first three are found by introducing small changes in the orientation of the celestial reference system around each of its principal axes, and deriving the corresponding changes in \mathbf{s} and \mathbf{a} (\mathbf{c} and \mathbf{g} being independent of the reference system); the last three are correspondingly found by introducing a small inertial spin of the reference system around each axis (see Sect. 6.1.5 for details). Introducing the $n \times 6$ matrix $\mathbf{V} = (\mathbf{v}_0, \dots, \mathbf{v}_5)$ we have

$$N\mathbf{V} = \mathbf{0}. \quad (36)$$

The singularity can in principle be removed by adding the six constraints $\mathbf{V}'\mathbf{x} = \mathbf{0}$, but in practice that is not necessary. It suffices to derive *one* particular solution $\tilde{\mathbf{x}}$ to the normal equations, then the whole solution space can be written $\tilde{\mathbf{x}} + \mathbf{V}\mathbf{z}$ for arbitrary $\mathbf{z} \in \mathbb{R}^6$. The vector \mathbf{z} is effectively determined by the frame rotator (Sect. 6.1), yielding the unique solution that best agrees with the adopted definition of the ICRS.

Quite apart from numerical rounding errors, it is not completely true that N is mathematically singular. Stellar aberration and parallax introduce some absolute knowledge about the reference system via the barycentric ephemeris of Gaia, which is expressed in ICRS and is not part of the adjustment process. However, since stellar aberration is at most 10^{-4} rad, the orientation error would have to be of the order of 1 mas for the

aberration effect to change by $0.1 \mu\text{as}$ (say). So, although absent in principle, the indeterminacy of the reference frame orientation and spin exists in practice. Since the orientation can be determined to much higher accuracy than 1 mas (by means of the optical counterparts of radio sources), the contribution of stellar aberration to the absolute frame knowledge can be neglected in practice. The same holds, a fortiori, for the much smaller parallax effect.

4.5. The simple iteration step

Consider the system

$$\mathbf{K}\mathbf{d} = \mathbf{b}, \quad \mathbf{K} = \begin{bmatrix} N_{ss} & \emptyset & \emptyset & \emptyset \\ N_{as} & N_{aa} & \emptyset & \emptyset \\ N_{cs} & N_{ca} & N_{cc} & \emptyset \\ N_{gs} & N_{ga} & N_{gc} & N_{gg} \end{bmatrix}. \quad (37)$$

where \mathbf{b} is the same right-hand side as in Eq. (30), and each \emptyset stands for a zero-filled sub-matrix of the appropriate dimensions. Although this system is of the same size as Eq. (30), it is fundamentally different in that it can be directly solved through a sequence of smaller systems,

$$\left. \begin{aligned} N_{ss}\mathbf{d}_s &= \mathbf{b}_s \\ N_{aa}\mathbf{d}_a &= \mathbf{b}_a - N_{as}\mathbf{d}_s \\ N_{cc}\mathbf{d}_c &= \mathbf{b}_c - N_{cs}\mathbf{d}_s - N_{ca}\mathbf{d}_a \\ N_{gg}\mathbf{d}_g &= \mathbf{b}_g - N_{gs}\mathbf{d}_s - N_{ga}\mathbf{d}_a - N_{gc}\mathbf{d}_c \end{aligned} \right\} \quad (38)$$

where each sub-system is of the form S, A, C, G described above, allowing it to be solved with relative ease. (Naturally, the resulting solution \mathbf{d} is also quite different from the \mathbf{x} in Eq. 30!) By means of Eqs. (29) and (34) the right-hand side in the second sub-system becomes

$$\begin{aligned} \mathbf{b}_a - N_{as}\mathbf{d}_s &= - \sum_l \frac{\partial R_l}{\partial \mathbf{a}} \left[R_l(\mathbf{s}^{\text{ref}}, \mathbf{a}^{\text{ref}}, \mathbf{c}^{\text{ref}}, \mathbf{g}^{\text{ref}}) + \frac{\partial R_l}{\partial \mathbf{s}'} \mathbf{d}_s \right] W_l \\ &= - \sum_l \frac{\partial R_l}{\partial \mathbf{a}} R_l(\mathbf{s}^{\text{ref}} + \mathbf{d}_s, \mathbf{a}^{\text{ref}}, \mathbf{c}^{\text{ref}}, \mathbf{g}^{\text{ref}}) W_l, \end{aligned} \quad (39)$$

where the linearity of $R_l(\mathbf{s}, \mathbf{a}, \mathbf{c}, \mathbf{g})$ has been used in a Taylor expansion around the reference values. This shows that the off-diagonal matrix N_{as} is not needed in order to compute the right-hand side of the second sub-system in Eq. (38), if only the residuals are computed *after* the source parameters have been updated by the solution of the first sub-system. Similarly, we find that the right-hand side in the third sub-system can be obtained from the residuals after updating both the source and attitude parameters, and so on. The off-diagonal sub-matrices in \mathbf{K} are therefore not needed, provided that the sub-vectors of unknowns are successively updated before the new right-hand sides are computed.

From the above it is clear that a single AGIS iteration, consisting of the successive application of the four update blocks S, A, C and G, is mathematically equivalent to an updating of the unknowns by $\mathbf{d} = \mathbf{K}^{-1}\mathbf{b}$. In the context of iterative solution algorithms, the matrix \mathbf{K} is referred to as the preconditioner of the normal equations system (Axelsson 1996).

As previously noted, we assume that the block-diagonal matrices N_{ss} , N_{aa} , N_{cc} , N_{gg} are all non-singular, and in fact positive definite, for a proper formulation of the S, A, C, and G blocks. This ensures that the preconditioner \mathbf{K} is non-singular, even though N is not.

In the course of the iterations, new right-hand sides of the normal equations will be computed, while the matrix remains

⁷ More precisely, N is properly semidefinite, so that $\mathbf{x}'N\mathbf{x} \geq 0$ for all $\mathbf{x} \neq \mathbf{0}$, with equality for some $\mathbf{x} \neq \mathbf{0}$.

essentially unchanged. From now on, let us express the residuals R_l not as functions of the parameter values s, a, c, g but in terms of the differential parameter vector \mathbf{x} relative to the reference values. The original right-hand side, denoted $\mathbf{b}^{(0)}$, corresponds to the initial differential parameter vector $\mathbf{x}^{(0)} = \mathbf{0}$. If \mathbf{x} (without a superscript) denotes the exact solution of Eq. (30), the initial error vector is $\mathbf{e}^{(0)} = \mathbf{x}^{(0)} - \mathbf{x} = -\mathbf{x}$. Although this vector is of course not known, we do know $-\mathbf{N}\mathbf{e}^{(0)} = \mathbf{N}\mathbf{x} = \mathbf{b}^{(0)}$. Solving the preconditioner system (37) gives $\mathbf{d}^{(0)} = \mathbf{K}^{-1}\mathbf{b}^{(0)}$ and the updated parameter vector $\mathbf{x}^{(1)} = \mathbf{x}^{(0)} + \mathbf{d}^{(0)}$. The new error vector $\mathbf{e}^{(1)} = \mathbf{x}^{(1)} - \mathbf{x}$ is again not known, but $-\mathbf{N}\mathbf{e}^{(1)} = \mathbf{b}^{(0)} - \mathbf{N}\mathbf{d}^{(0)} = \mathbf{b}^{(1)}$ is obtained by inserting the updated parameters in the right-hand side of Eq. (29), by the same argument as in Eq. (39). Generalizing, we have the so-called *simple iteration scheme*

$$\left. \begin{aligned} \mathbf{b}^{(k)} &= - \sum_l \frac{\partial R_l}{\partial \mathbf{x}} R_l(\mathbf{x}^{(k)}) W_l \\ \mathbf{d}^{(k)} &= \mathbf{K}^{-1} \mathbf{b}^{(k)} \\ \mathbf{x}^{(k+1)} &= \mathbf{x}^{(k)} + \mathbf{d}^{(k)} \end{aligned} \right\} \quad k = 0, 1, \dots \quad (40)$$

For the successive right-hand sides we find by recursion

$$\mathbf{b}^{(k+1)} = \mathbf{b}^{(k)} - \mathbf{N}\mathbf{d}^{(k)} = (\mathbf{I} - \mathbf{N}\mathbf{K}^{-1})\mathbf{b}^{(k)} = \tilde{\mathbf{B}}^{k+1}\mathbf{b}^{(0)}, \quad (41)$$

where

$$\tilde{\mathbf{B}} = \mathbf{I} - \mathbf{N}\mathbf{K}^{-1}. \quad (42)$$

For the successive updates and errors we find

$$\mathbf{d}^{(k+1)} = \mathbf{K}^{-1}\mathbf{b}^{(k+1)} = (\mathbf{I} - \mathbf{K}^{-1}\mathbf{N})\mathbf{d}^{(k)} = \mathbf{B}^{k+1}\mathbf{d}^{(0)}, \quad (43)$$

$$\mathbf{e}^{(k+1)} = \mathbf{e}^{(k)} + \mathbf{d}^{(k)} = (\mathbf{I} - \mathbf{K}^{-1}\mathbf{N})\mathbf{e}^{(k)} = \mathbf{B}^{k+1}\mathbf{e}^{(0)}. \quad (44)$$

where

$$\mathbf{B} = \mathbf{I} - \mathbf{K}^{-1}\mathbf{N} \quad (45)$$

is called the iteration matrix (Axelsson 1996). Equations (41)–(45) are of great theoretical interest, as explained below, although none of them is used in the actual computations.

The convergence behaviour of the simple iteration scheme can largely be understood by means of these relations and especially in terms of the properties of the iteration matrix \mathbf{B} governing the sequence of updates \mathbf{d} and errors \mathbf{e} , and the adjunct matrix $\tilde{\mathbf{B}}$ governing the sequence of right-hand sides. It is well known (e.g., Axelsson 1996) that the simple iteration scheme in Eq. (40) converges (that is $\mathbf{x}^{(k)} \rightarrow \mathbf{x}$) for arbitrary starting approximation if and only if $\rho(\mathbf{B}) < 1$. Here, $\rho(\mathbf{B})$ is the spectral radius of \mathbf{B} , i.e., the largest absolute value of its eigenvalues. Under this condition we have $\mathbf{d}^{(k)} \rightarrow \mathbf{0}$ and $\mathbf{e}^{(k)} \rightarrow \mathbf{0}$ for $k \rightarrow \infty$. Also the right-hand side $\mathbf{b}^{(k)} = \mathbf{K}\mathbf{d}^{(k)} \rightarrow \mathbf{0}$ under the same condition.⁸

As discussed in Sect. 4.4, the normal matrix \mathbf{N} is singular and its null space spanned by the $n \times 6$ matrix \mathbf{V} . Therefore,

$$\mathbf{B}\mathbf{V} = \mathbf{V} - \mathbf{K}^{-1}\mathbf{N}\mathbf{V} = \mathbf{V} \quad (46)$$

which shows that \mathbf{B} has a six-fold eigenvalue equal to 1, with the corresponding eigenvectors spanning $\mathcal{N}(\mathbf{N})$. The spectral radius of \mathbf{B} is therefore not less than 1, and the simple iteration scheme does not converge for arbitrary initial errors.

⁸ \mathbf{B} and $\tilde{\mathbf{B}} = \mathbf{K}\mathbf{B}\mathbf{K}^{-1}$ have the same spectral radius; indeed, their eigenvalues are the same, as can be seen from the characteristic polynomial $\det(z\mathbf{I} - \mathbf{B}) = \det(\mathbf{K})(z\mathbf{I} - \mathbf{B})\det(\mathbf{K}^{-1}) = \det(z\mathbf{I} - \mathbf{K}\mathbf{B}\mathbf{K}^{-1})$ being the same for the two matrices (A. Bombrun, private communication).

This is not a real problem, for the following reason. First, let us note that \mathbf{B} is not a full-rank matrix. Indeed, a direct calculation of Eq. (45) using the expression for \mathbf{K} in Eq. (37) shows that the first n_s columns of \mathbf{B} are zero.⁹ Thus (at least) n_s of its eigenvalues are identically zero. A corresponding set of orthogonal unit vectors is given by the columns of the $n \times n_s$ matrix

$$\mathbf{Z} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \quad \left. \begin{array}{l} \} n_s \text{ rows} \\ \} n - n_s \text{ rows} \end{array} \right\}, \quad (47)$$

so that $\mathbf{B}\mathbf{Z} = \mathbf{0}$. We assume that the remaining $n - n_s - 6$ eigenvalues satisfy $0 < |\lambda| < 1$. Thus, if $\mathbf{\Lambda}$ is the diagonal matrix containing these eigenvalues and \mathbf{U} a matrix of size $n \times (n - n_s - 6)$ whose columns are made up of the corresponding eigenvectors, we have $\mathbf{B}\mathbf{U} = \mathbf{U}\mathbf{\Lambda}$. The columns of \mathbf{Z} , \mathbf{U} and \mathbf{V} together span \mathbb{R}^n , and it is therefore possible to decompose the solution vector as

$$\mathbf{x} = \mathbf{Z}\mathbf{x}_s + \mathbf{U}\mathbf{y} + \mathbf{V}\mathbf{z} \quad (48)$$

where $\mathbf{x}_s \in \mathbb{R}^{n_s}$ is the ‘source’ part of \mathbf{x} , $\mathbf{y} \in \mathbb{R}^{n-n_s-6}$ and $\mathbf{z} \in \mathbb{R}^6$. Since $-\mathbf{e}^{(0)} = \mathbf{x}$ we find

$$-\mathbf{e}^{(1)} = -\mathbf{B}\mathbf{e}^{(0)} = \mathbf{B}\mathbf{Z}\mathbf{x}_s + \mathbf{B}\mathbf{U}\mathbf{y} + \mathbf{B}\mathbf{V}\mathbf{z} = \mathbf{U}\mathbf{\Lambda}\mathbf{y} + \mathbf{V}\mathbf{z}, \quad (49)$$

and by recursion

$$-\mathbf{e}^{(k)} = \mathbf{U}\mathbf{\Lambda}^k\mathbf{y} + \mathbf{V}\mathbf{z}. \quad (50)$$

The first term clearly vanishes for $k \rightarrow \infty$ if $\rho(\mathbf{\Lambda}) < 1$, as we have assumed. The second term, which is the projection of \mathbf{x} on the null space of \mathbf{N} , remains unchanged by the iterations. The update vector can be written

$$\mathbf{d}^{(k)} = \mathbf{e}^{(k+1)} - \mathbf{e}^{(k)} = \mathbf{U}\mathbf{\Lambda}^k(\mathbf{I} - \mathbf{\Lambda})\mathbf{y}, \quad (51)$$

which vanishes under the same condition. The same is then true for the right-hand side $\mathbf{b}^{(k)} = \mathbf{K}\mathbf{d}^{(k)}$. Effectively, this means that we can ignore the singularity of \mathbf{N} in the simple iteration scheme; it will converge to some valid solution of the (singular) normal equations, and the convergence process can be monitored by means of the vectors $\mathbf{d}^{(k)}$ and $\mathbf{b}^{(k)}$. After convergence, the required null-space components can be found and added by means of the frame rotator.

To summarize, we have shown that the simple iteration scheme converges in the desired way, provided that the spectral radius of \mathbf{B} , not counting the eigenvalues related to the null space of \mathbf{N} , is less than 1. It is well known (e.g., Golub & van Loan 1996) that this condition is satisfied for any symmetric positive definite \mathbf{N} , using the Gauss–Seidel preconditioner. However, the spectral radius may be very close to 1, implying very slow convergence. In the case of AGIS the situation is more complex, and convergence is in practice demonstrated through simulations (Sect. 7), but the theoretical background outlined above is of great help when interpreting the results.

⁹ That the first n_s columns in the iteration matrix are zero means that the results of the next iteration are independent of the current source parameters $\mathbf{x}_s^{(k)}$. This may seem surprising at first, but is a simple consequence of the fact that each iteration starts with the source update block (S). In this block, the updated source parameters depend on the previous attitude, calibration and global parameters, but not on the previous source parameters.

4.6. Order of the block processes

In the preceding sections we have assumed that the four blocks are always executed in the sequence SACG, and that the updated parameters resulting from each block is used in the subsequent blocks. The particular preconditioner \mathbf{K} in Eq. (37) incorporates these assumptions. We will now discuss variants of this scheme and show that they can be mathematically represented by different preconditioners.

Each AGIS iteration always starts with the source update block (S). The main reason for starting with S rather than A, for example, is that the observations can be inspected and analysed one source at a time in order to estimate the quality of the data, identify possible outliers, and check whether the source is suitable as a primary source. The identification of outliers, in particular, requires several ‘inner’ iterations of the source observation equations (Sect. 5.1.2), which can be done with relative ease because of the small number of data points involved.

The order of the remaining blocks ACG adopted in the preceding sections is more or less arbitrary. It is easy to see through an analogy with Eqs. (37)–(38) that the sequence SCAG (for example) is mathematically represented by the alternative preconditioner

$$\mathbf{K} = \begin{bmatrix} N_{ss} & \emptyset & \emptyset & \emptyset \\ N_{as} & N_{aa} & N_{ac} & \emptyset \\ N_{cs} & \emptyset & N_{cc} & \emptyset \\ N_{gs} & N_{ga} & N_{gc} & N_{gg} \end{bmatrix}, \quad (52)$$

and that other permutations of the sequence are similarly obtained by transposing the corresponding off-diagonal blocks. Changing the preconditioner means changing the iteration matrix (45) and therefore possibly also its eigenvalues, which in turn govern the rate of convergence.

From a mathematical viewpoint, the choice of the starting block (S in our case) determines the initial update $\mathbf{d}^{(0)}$ (cf. footnote 9) and therefore influences all subsequent updates. However, this particular choice is not expected to have much influence on the convergence rate after a number of iterations, since the S block has no special status in the periodically repeated sequence ...SACGSACGSACGSA.... Thus we may surmise that the different iteration matrices for the cyclically permuted sequences SACG, ACGS, CGSA and GSAC do in fact have the same set of eigenvalues. For symmetry reasons it can also be surmised that the reversed sequences GCAS, SGCA, ASGC and CASG have the same eigenvalues as the original sequences. Thus, there are in fact only three sets of sequences with possibly distinct convergence behaviour, represented by SACG (or SGCA), SAGC (or SCGA), and SCAG (or SGAC) if we take S to be the starting block. There is no obvious way of knowing a priori which of the three possibilities is to be preferred, or even if they are significantly different.

Apart from the various permutations discussed above, there is another way to modify the AGIS iteration scheme, which can also be described in terms of the preconditioner. This modification is related to the practical organization of the flow of data to the different block processes. The current implementation of the simple iteration scheme differs somewhat from the description in Sect. 4.5, and the block updates are in fact practically organized as follows:

1. Initialize the iteration counter, $k = 0$
2. Choose starting values for all unknowns: $\mathbf{s}^{(0)}$, $\mathbf{a}^{(0)}$, $\mathbf{c}^{(0)}$, $\mathbf{g}^{(0)}$.
3. Estimate $\mathbf{s}^{(k+1)}$ using $\mathbf{a}^{(k)}$, $\mathbf{c}^{(k)}$, $\mathbf{g}^{(k)}$.
4. Estimate $\mathbf{a}^{(k+1)}$ using $\mathbf{s}^{(k+1)}$, $\mathbf{c}^{(k)}$, $\mathbf{g}^{(k)}$.
5. Estimate $\mathbf{c}^{(k+1)}$ using $\mathbf{s}^{(k+1)}$, $\mathbf{a}^{(k+1)}$, $\mathbf{g}^{(k)}$.

6. Estimate $\mathbf{g}^{(k+1)}$ using $\mathbf{s}^{(k+1)}$, $\mathbf{a}^{(k)}$, $\mathbf{c}^{(k)}$.
7. Increment k and go to Step 3.

The crucial difference compared with Eq. (38) is that the C block in Step 5 does not use the updated attitude but the old one, and that the G block in Step 6 similarly uses the ‘old’ attitude and calibration parameters. This has the practical advantage that the A, C and G blocks can be carried out in parallel, with big savings in terms of the amounts of data that have to be exchanged between the processes (see Sect. 7). Schematically, this can be represented as S[ACG], where the blocks in brackets are (or can be) executed in parallel (and so the order of the bracketed blocks does not matter). The corresponding preconditioner is

$$\mathbf{K} = \begin{bmatrix} N_{ss} & \emptyset & \emptyset & \emptyset \\ N_{as} & N_{aa} & \emptyset & \emptyset \\ N_{cs} & \emptyset & N_{cc} & \emptyset \\ N_{gs} & \emptyset & \emptyset & N_{gg} \end{bmatrix}. \quad (53)$$

Yet other variants may be considered, for example S[AC]G, where Step 6 uses the updated attitude and calibration parameters, with preconditioner

$$\mathbf{K} = \begin{bmatrix} N_{ss} & \emptyset & \emptyset & \emptyset \\ N_{as} & N_{aa} & \emptyset & \emptyset \\ N_{cs} & \emptyset & N_{cc} & \emptyset \\ N_{gs} & N_{ga} & N_{gc} & N_{gg} \end{bmatrix}, \quad (54)$$

and [SACG], for which

$$\mathbf{K} = \begin{bmatrix} N_{ss} & \emptyset & \emptyset & \emptyset \\ \emptyset & N_{aa} & \emptyset & \emptyset \\ \emptyset & \emptyset & N_{cc} & \emptyset \\ \emptyset & \emptyset & \emptyset & N_{gg} \end{bmatrix}. \quad (55)$$

This last case is known as the (block) Jacobi method, while the use of a full triangular preconditioner as in Eq. (37) is known as the (block) Gauss–Seidel method (Axelsson 1996). As we have seen, the currently implemented simple iteration scheme is intermediate between the Jacobi and Gauss–Seidel methods.

Intuitively, we expect the Gauss–Seidel method to converge more quickly than the Jacobi or any intermediate method, simply because each block then uses the most recent (and, presumably, best) estimates of the parameters. However, our practical experience with AGIS shows that the ‘difficult’ part of the problem is to disentangle the source and attitude parameters. For example, the calibration parameters are generally found to converge much faster than the source and attitude parameters. Thus it does not seem to matter much if Step 5 above uses $\mathbf{a}^{(k+1)}$ or $\mathbf{a}^{(k)}$, i.e., whether the sub-matrix N_{ca} is included or not in the preconditioner. A similar argument can be made concerning the G block, provided some measures are taken to decorrelate the global parameters from the source parameters (Sect. 5.4). Thus, the various intermediate methods are probably nearly as good as the Gauss–Seidel method, in terms of the convergence rate, and the precise scheme may then rather be determined by practical considerations. With the present data processing architecture, the favoured scheme is S[ACG] as described above.

4.7. Accelerated iteration, conjugate gradients and the hybrid iteration scheme

The ‘simple iteration’ (SI) scheme described above was the starting point for a long development towards a fully functional scheme with much improved convergence properties. The main

stages in this development were the ‘accelerated simple iteration’ (ASI), the conjugate gradients (CG), and finally the fully flexible ‘hybrid scheme’ (A/SI-CG) to be used in the final implementation of AGIS. As much of this development has at most historical interest, only a brief outline is given here.

Already in the very early implementation of the simple iteration scheme it was observed that convergence was slower than (naively) expected, and that after some iterations, the updates always seemed to go in the same direction, forming a geometrically (exponentially) decreasing series. With the hindsight of the analysis in Sect. 4.5 this behaviour is very easily understood: the persistent pattern of updates is roughly proportional to the eigenvector of the largest eigenvalue of the iteration matrix, and the (nearly constant) ratio of the sizes of successive updates is the corresponding eigenvalue. From this realization it was natural to test an acceleration method based on a Richardson-type extrapolation of the updates. The idea is simply that if the updates in two successive iterations are roughly proportional to each other, $\mathbf{d}^{(k+1)} \simeq \lambda \mathbf{d}^{(k)}$, with $|\lambda| < 1$, then we can infer that the next update is again a factor λ smaller than $\mathbf{d}^{(k+1)}$, and so on. The sum of all the updates after iteration k can therefore be estimated as $\mathbf{d}^{(k+1)} + \lambda \mathbf{d}^{(k+1)} + \lambda^2 \mathbf{d}^{(k+1)} + \dots = (1 - \lambda)^{-1} \mathbf{d}^{(k+1)}$. Thus, in iteration $k + 1$ we apply an acceleration factor $1/(1 - \lambda)$ based on the current estimate of the ratio λ . This accelerated simple iteration (ASI) scheme is seen to be a variant of the well-known successive overrelaxation method (Axelsson 1996). The factor λ is estimated by statistical analysis of the parallax updates for a small fraction of the sources; the parallax updates are used for this analysis, since they are unaffected by a possible change in the frame orientation between successive iterations. With this simple device, the number of iterations for full convergence was reduced roughly by a factor 2.

Both the simple iteration and the accelerated simple iteration belongs to a much more general class of solution methods known as Krylov subspace approximations. The sequence of updates $\mathbf{d}^{(k)}$, $k = 0 \dots K - 1$ generated by the first K simple iterations constitute the basis for the K -dimensional subspace of the solution space, known as the Krylov subspace for the given matrix and right-hand side (e.g., Greenbaum 1997; van der Vorst 2003). Krylov methods compute approximations that, in the k th iteration, belongs to the k -dimensional Krylov subspace. But whereas the simple and accelerated iteration schemes, in the k th iteration, use updates that are just proportional to the k th basis vector, more efficient algorithms generate approximations that are (in some sense) optimal linear combinations of all k basis vectors. Conjugate gradients (CG) is one of the best-known such methods, and possibly the most efficient one for general symmetric positive-definite matrices. (e.g., Axelsson 1996; Björck 1996; van der Vorst 2003). Its implementation within the AGIS framework is more complicated, but has been considered in detail by Bombrun et al. (2011). As it provides significant advantages over the SI and ASI schemes in terms of convergence speed, this algorithm has been chosen as the baseline method for the astrometric core solution of Gaia (see below however). From practical experience, we have found that CG is roughly a factor 2 faster than ASI, or a factor 4 faster than the SI scheme. Like SI, the CG algorithm uses a preconditioner and can be formulated in terms of the S, A, C and G blocks, so the subsequent description of these blocks remains valid. In the terminology of Bombrun et al. (2011) the process of solving the preconditioner system $\mathbf{K}\mathbf{d} = \mathbf{b}$ is the *kernel* operation common to all these solution methods, which only differ in how the updates are applied according to the various *iteration schemes*.

The CG algorithm assumes that the normal matrix is constant in the course of the iterations. This is not strictly true if the observation weights are allowed to change as functions of the residuals, as will be required for efficient outlier elimination (Sect. 5.1.2). Using the CG algorithm together with the weight-adjustment scheme described below could therefore lead to instabilities, i.e., a reduced convergence rate or even non-convergence. On the other hand, the SI scheme is extremely stable with respect to all such modifications in the course of the iterations, as can be expected from the interpretation of the SI scheme as the successive and independent application of the different solution blocks. The finally adopted algorithm is therefore a *hybrid scheme* combining SI (or ASI) and CG, where SI is used initially, until the weights have settled, after which CG is turned on. A temporary switch back to SI, with an optional re-adjustment of the weights, may be employed after a certain number of CG iterations; this could avoid some problems due to the accumulation of numerical rounding errors in CG.

5. Updating processes

In this section we describe in some detail each of the updating blocks S, A, C and G that form the basis (or kernel process) for the AGIS iteration loop.

5.1. Source updating (S)

5.1.1. The normal equations

The astrometric model for the sources is given in Sect. 3.2. In the source update block (S) the source parameters \mathbf{s} are improved by solving the first line in Eq. (38). According to Eqs. (29) and (32) this can be done for one source (i) at a time by solving the following system of equations for the update \mathbf{d}_i of the five astrometric parameters in \mathbf{s}_i :

$$\left[\sum_{l \in i} \frac{\partial R_l}{\partial \mathbf{s}_i} \frac{\partial R_l}{\partial \mathbf{s}_i'} \mathbf{W}_l \right] \mathbf{d}_i = - \sum_{l \in i} \frac{\partial R_l}{\partial \mathbf{s}_i} R_l(\mathbf{s}_i) \mathbf{W}_l. \quad (56)$$

Here $\mathbf{W}_l = w_l / (\sigma_l^2 + \epsilon_l^2)$, with σ_l denoting the given formal standard uncertainty of observation l , expressed as an angle. w_l and ϵ_l are, respectively, the downweighting factor and excess noise introduced in Sect. 3.6. In Eq. (56) the dependence of R_l on \mathbf{a} , \mathbf{c} and \mathbf{g} has been suppressed, since the system is solved with these parameters fixed.

In matrix notation, the normal equations (56) can be written

$$\mathbf{A}_i' \mathbf{W}_i \mathbf{A}_i \mathbf{d}_i = \mathbf{A}_i' \mathbf{W}_i \mathbf{h}_i, \quad (57)$$

where

$$\mathbf{A}_i = \left[- \frac{\partial R_l}{\partial \mathbf{s}_i'} \right]_{l \in i} \quad (58)$$

is an $n_i \times 5$ matrix with $n_i = \sum_{l \in i} 1$ the number of observations of source i (typically $n_i \sim 800$),

$$\mathbf{h}_i = [R_l(\mathbf{s}_i)]_{l \in i} \quad (59)$$

is a column matrix of length n_i , and \mathbf{W}_i is a diagonal matrix containing the weight factors \mathbf{W}_l . Although the residuals R_l are in principle non-linear functions of \mathbf{s}_i , this non-linearity can be neglected if the parameters are close enough to the final solution, i.e., if the resulting update \mathbf{d}_i is small enough. The partial derivatives in \mathbf{A}_i can then be regarded as fixed throughout the updating process, and \mathbf{A}_i and \mathbf{h}_i can immediately be computed when entering the source updating. The weight matrix \mathbf{W}_i , on the other

hand, depends on w_l and ϵ_l , which are modified as part of the process as explained below.

For given \mathbf{W}_i the system of equations (57) is solved using the Cholesky algorithm (Appendix C). At the end of the source updating, the full (5×5) inverse matrix is computed, providing a first-order estimate of the covariance of the astrometric parameters in \mathbf{s}_i (see Sect. 6.3).

We note that Eq. (57) corresponds to the overdetermined system of observation equations

$$\mathbf{A}_i \mathbf{d}_i \approx \mathbf{h}_i \quad (\text{weight matrix } \mathbf{W}_i). \quad (60)$$

After solution, the updated residuals are contained in the vector $\mathbf{R}_i = \mathbf{h}_i - \mathbf{A}_i \mathbf{d}_i$, and the contribution of the source to the objective function Q is given by

$$Q_i = \mathbf{R}_i' \mathbf{W}_i \mathbf{R}_i. \quad (61)$$

5.1.2. The inner iteration: identifying outliers and estimating the excess source noise

Since the weight matrix \mathbf{W}_i depends on the downweighting factors w_l and the excess noises ϵ_l , which in turn depend on the updated residuals \mathbf{R}_i , the normal equations (57) are non-linear and must be solved by iteration. We refer to this as the *inner iteration* of the source update, to distinguish it from the AGIS iteration discussed in Sect. 4.5.

Using Eq. (28) we can write

$$W_l = \frac{w_l}{\sigma_l^2 + \epsilon_a^2(t_l) + \epsilon_i^2} = \frac{w_l}{\tilde{\sigma}_l^2 + \epsilon_i^2}, \quad (62)$$

where $\tilde{\sigma}_l = [\sigma_l^2 + \epsilon_a^2(t_l)]^{1/2}$ is the formal standard uncertainty of the observation adjusted for the excess attitude noise, which, when entering the source update, is assumed to be known from a previous attitude update (Sect. 5.2.5). In the very first source update the excess attitude noise must be set to the mean errors of the pre-AGIS attitude parameter estimates which are used for starting up the iterations. The downweighting factors depend on the normalized residuals, i.e.,

$$w_l = w \left(\frac{R_l}{\sqrt{\tilde{\sigma}_l^2 + \epsilon_i^2}} \right), \quad (63)$$

where the function $w(z)$ is such that $w(z) = 1$ for $|z| \lesssim 3$ and gradually decreasing to 0 for larger $|z|$ (see Eq. 66).

The inner iteration actually consists of two nested procedures: an outer one which determines the downweighting factors w_l , and an inner one which determines the excess source noise ϵ_i for a fixed set of w_l . We begin by considering the inner procedure.

For fixed downweighting factors, the source update aims to minimize Q_i in Eq. (61) with respect to the unknown \mathbf{d}_i and ϵ_i . But it is immediately seen that Q_i can be made arbitrarily small just by making ϵ_i large enough. Consequently, we cannot use unconstrained minimization to solve this problem. The necessary constraint is provided by the condition that Q_i , under the assumption that the excess noises have been correctly estimated and the outliers properly removed, should follow the chi-square distribution with $\nu = n_i - n_{\text{out}} - 5$ degrees of freedom. Here n_i is the number of observations of source i , n_{out} the number of outliers and 5 the number of astrometric parameters estimated. In particular, the expected value is $E(Q_i) = \nu$. The number of outliers n_{out} is estimated by counting the number of observations

with $w_l < 0.2$. This limit was empirically found to give a reasonable estimate of the actual number of outliers in a variety of simulated cases.

For given ϵ_i the weight matrix is now known, Eq. (57) can be trivially solved and the residuals \mathbf{R}_i computed. We thus define the function $Q(\epsilon_i^2) = \mathbf{R}_i' \mathbf{W}_i \mathbf{R}_i$. The excess source variance is then taken to be

$$\epsilon_i^2 = \begin{cases} 0 & \text{if } Q(0) \leq \nu, \\ \text{solution of } Q(\epsilon_i^2) = \nu & \text{otherwise.} \end{cases} \quad (64)$$

In the second case, the non-linear equation $Q(y) = \nu$ is iteratively solved by a series of improvements $\Delta y = (1 - Q(y)/\nu)Q(y)/Q'(y)$, starting from $y = 0$. Typically, 2–3 iterations are sufficient.¹⁰

This procedure returns a positive ϵ_i as soon as $Q(0) > \nu$. If the resulting ϵ_i is much smaller than the typical σ_l of the source, it is probably not significant. The significance of the ϵ_i can more easily be judged from an auxiliary statistic that can be computed almost for free. Under the null hypothesis ($\epsilon_i = 0$) we know that $Q(0) \sim \chi_\nu^2$, so the expected value is ν and the variance 2ν . Thus we may take

$$D = \frac{Q(0) - \nu}{\sqrt{2\nu}} \quad (65)$$

as a measure of the significance of the estimated ϵ_i , with $D > 2$ indicating a probably significant value.

Having determined ϵ_i , and a corresponding set of residuals \mathbf{R}_i , the downweighting factors can immediately be computed from Eq. (63). However, having changed the downweighting factors (and possibly ν) it is now necessary to repeat the estimation of ϵ_i and \mathbf{d}_i with the new set w_l . Typically, four such iterations are found to be sufficient. The only remaining problem is how to start the iterations, that is the initial selection of the downweights w_l . It is not possible to start by assuming $w_l = 1$ for all the observations, since we must take into account that some small fraction of the data could be utterly wrong. Such gross outliers, if not removed already from the start, would severely slow down or even prevent the convergence of the inner iterations. The adopted solution is to make an extremely robust estimation of the standard deviation of the initial residuals (contained in \mathbf{h}_i), from which the initial downweightings are obtained. This robust standard deviation ς_i is calculated as half the intersextile range of the elements in \mathbf{h}_i , whereupon the initial $w_l = w(h_l/\varsigma_i)$.

After convergence of the inner iteration, the statistical weight of the source W_i is computed according to Eq. (119). This quantity, together with ϵ_i and the magnitude, are the most important indicators for the selection of primary sources (Sect. 6.2.2).

The weight function $w(z)$ currently used is the following:

$$w(z) = \begin{cases} 1 & \text{if } |z| \leq 2 \\ 1 - 1.773735t^2 + 1.141615t^3 & \text{if } 2 \leq |z| < 3 \\ \exp(-|z|/3) & \text{otherwise,} \end{cases} \quad (66)$$

where $t = |z| - 2$ and the numerical constants have been chosen to make a smooth transition at $|z| = 3$. The exponential decay for $|z| > 3$ provides a dramatic weight reduction for large residuals; e.g., at 10 sigmas we have $w(10) \approx 0.036$, while at 100 sigmas we have $w(100) \approx 3 \times 10^{-15}$.

¹⁰ This iteration formula can be derived by matching the rational approximation $Q(y) \approx a/(b+y)$ to the value and derivative of $Q(y)$ at the current point y . Compared with the standard Newton–Raphson method, which uses a linear approximation around the current y , the present formula converges much quicker due to the more reasonable behaviour of the rational approximation especially for large y .

Among the many different weight functions proposed in the literature, the so-called Huber estimator (Huber 1981) using

$$w_H(z) = \begin{cases} 1 & \text{if } |z| \leq c \\ c(2|z| - c)z^{-2} & \text{otherwise} \end{cases} \quad (67)$$

has been considered as an alternative to Eq. (66), e.g., with $c = 2$. This gives a much slower weight decay for large residuals, e.g., $w_H(10) = 0.36$ and $w_H(100) = 0.0396$. Future experiments may decide which weight function will finally be used for AGIS.

5.1.3. Calculation of partial derivatives

The calculation of the partial derivatives in Eq. (58) is done as follows. From Eqs. (25) and (26) we have, by means of (13),

$$-\frac{\partial R_l^{\text{AL}}}{\partial s'_i} = \frac{\partial \varphi_l}{\partial s'_i}, \quad -\frac{\partial R_l^{\text{AC}}}{\partial s'_i} = \frac{\partial \zeta_l}{\partial s'_i}. \quad (68)$$

In analogy with Eq. (5) we introduce the auxiliary vectors $\mathbf{m}_l = \langle \mathbf{z} \times \mathbf{u}_l \rangle$ and $\mathbf{n}_l = \mathbf{u}_l \times \mathbf{m}_l$, which together with \mathbf{u}_l form the normal triad $[\mathbf{m}_l \ \mathbf{n}_l \ \mathbf{u}_l]$ with respect to the SRS; its components in the SRS are given by the columns of the matrix

$$\mathbf{S}'[\mathbf{m}_l \ \mathbf{n}_l \ \mathbf{u}_l] = \begin{bmatrix} -\sin \varphi_l & -\sin \zeta_l \cos \varphi_l & \cos \zeta_l \cos \varphi_l \\ \cos \varphi_l & -\sin \zeta_l \sin \varphi_l & \cos \zeta_l \sin \varphi_l \\ 0 & \cos \zeta_l & \sin \zeta_l \end{bmatrix}. \quad (69)$$

By differentiation of the last column we obtain

$$\frac{\partial \mathbf{u}_l}{\partial s'_i} = \mathbf{m}_l \frac{\partial \varphi_l}{\partial s'_i} \cos \zeta_l + \mathbf{n}_l \frac{\partial \zeta_l}{\partial s'_i}, \quad (70)$$

and thus

$$-\frac{\partial R_l^{\text{AL}}}{\partial s'_i} = \mathbf{m}'_l \frac{\partial \mathbf{u}_l}{\partial s'_i} \sec \zeta_l, \quad -\frac{\partial R_l^{\text{AC}}}{\partial s'_i} = \mathbf{n}'_l \frac{\partial \mathbf{u}_l}{\partial s'_i}. \quad (71)$$

These expressions can be evaluated in any coordinate system, but perhaps most conveniently in the SRS using $\mathbf{S}'\mathbf{m}_l$ and $\mathbf{S}'\mathbf{n}_l$ from Eq. (69), and

$$\left\{ \frac{\partial (\mathbf{S}'\mathbf{u}_l)}{\partial s'_i}, 0 \right\} = \mathbf{q}_l^{-1} \left\{ \frac{\partial (\mathbf{C}'\mathbf{u}_l)}{\partial s'_i}, 0 \right\} \mathbf{q}_l \quad (72)$$

from Eq. (11). To compute the partial derivatives of $\mathbf{C}'\mathbf{u}_l$, we first obtain from Eq. (4) the derivatives¹¹ of the coordinate direction as:

$$\begin{aligned} \frac{\partial \bar{\mathbf{u}}_l}{\partial \alpha_{*i}} &= \mathbf{p}_i, & \frac{\partial \bar{\mathbf{u}}_l}{\partial \delta_i} &= \mathbf{q}_i, & \frac{\partial \bar{\mathbf{u}}_l}{\partial \varpi_i} &= -(\mathbf{I} - \mathbf{r}_i \mathbf{r}'_i) \mathbf{b}_G(t_l) / A_u, \\ \frac{\partial \bar{\mathbf{u}}_l}{\partial \mu_{\alpha * i}} &= \mathbf{p}_i \tau_l, & \frac{\partial \bar{\mathbf{u}}_l}{\partial \mu_{\delta i}} &= \mathbf{q}_i \tau_l, \\ \frac{\partial \bar{\mathbf{u}}_l}{\partial \mu_{r i}} &= (\mathbf{I} - \mathbf{r}_i \mathbf{r}'_i) \mathbf{b}_G(t_l) \varpi_i \tau_l / A_u - (\mathbf{p}_i \mu_{\alpha * i} + \mathbf{q}_i \mu_{\delta i}) \tau_l^2, \end{aligned} \quad (73)$$

where $\tau_l = t_{Bl} - t_{ep}$ for brevity, and we have used the normal triad in the celestial reference system, Eq. (5). Although usually only the first five derivatives are needed (see however Sect. 6.3), the last equation gives, for completeness, the derivative with respect to the sixth astrometric parameter $\mu_{r i}$: the first term corresponds

to the secular change in parallax and the second to the perspective acceleration.

The rigorous transformation from $\bar{\mathbf{u}}_l$ to \mathbf{u}_l is quite complex, but by far the largest difference ($\sim 10^{-4}$ rad) is due to stellar aberration. By comparison, gravitational light deflection by the Sun is $\sim 2 \times 10^{-8}$ rad. While the rigorous transformation is required to compute the vector \mathbf{u}_l itself, some simplifications can be accepted when computing the partial derivatives. Indeed, for this purpose it is sufficient to consider the classical stellar aberration formula,

$$\mathbf{u}_l \simeq \langle \bar{\mathbf{u}}_l + \mathbf{v}_G(t_l) c^{-1} \rangle, \quad (74)$$

accurate to first order in v_G/c , where $\mathbf{v}_G = d\mathbf{b}_G/dt$ is the barycentric coordinate velocity of Gaia and c the speed of light. To a relative precision better than 10^{-6} we then have

$$\frac{\partial \mathbf{u}_l}{\partial s'_i} \simeq \left[\left(1 - \frac{\bar{\mathbf{u}}_l \mathbf{v}_G}{c} \right) \mathbf{I} - \frac{\bar{\mathbf{u}}_l \mathbf{v}'_G}{c} \right] \frac{\partial \bar{\mathbf{u}}_l}{\partial s'_i}, \quad (75)$$

where \mathbf{I} is the 3×3 identity matrix.

5.2. Attitude updating (A)

5.2.1. The normal equations

The attitude model using B-splines to represent the components of the attitude quaternion as functions of time is described in Sect. 3.3. (For reference purposes, conventions for notation and some important properties of splines and B-splines are explained in Appendix B.1.) In the attitude update process (A) the attitude parameters \mathbf{a} are improved by solving the second sub-system in Eq. (38). Recalling that \mathbf{a} is divided into sub-vectors of length 4, representing the quaternions \mathbf{a}_n in Eq. (10), the n th set of four equations can be written, using Eqs. (29) and (33),

$$\sum_{m=n-M+1}^{n+M-1} \left[\sum_{l \in L_n \cap L_m} \frac{\partial R_l}{\partial \mathbf{a}_n} \frac{\partial R_l}{\partial \mathbf{a}'_m} W_l \right] \mathbf{d}_m = - \sum_{l \in L_n} \frac{\partial R_l}{\partial \mathbf{a}_n} R_l(s, \mathbf{a}, c, g) W_l. \quad (76)$$

Here L_n stands for the set of observations occurring within the support of $B_n(t)$, i.e., $L_n = \{l | \tau_n \leq t_l < \tau_{n+M}\}$, where M is the order of the spline. On the right-hand side, it is understood that the residuals R_l are calculated for the most recent source parameters s (i.e., from the preceding source updating), while the attitude, calibration and global parameters are the not-yet-updated ones, as explained in Sect. 4.6. The weights W_l are the ones calculated in the source updating.

The structure of the symmetric matrix N_{aa} is shown in Fig. 5. If each quaternion component is represented by a spline of order M with N degrees of freedom (so $n = 0 \dots N-1$), the total number of unknowns is $4N$ and the average bandwidth of N_{aa} (counting non-zero elements from the diagonal up) is $4(M-1) + 2.5$. Including the right-hand side, the total number of reals that need to be stored for the normal equations is therefore $\simeq (16M-2)N$ or $62N$ for cubic splines. With a knot interval of about 15 s, about 3 MB is required to store the attitude normal equations for one day of observations. Thus it is completely realistic to store the attitude normal equations for the entire mission in primary memory. Cholesky factorization of the normal equations does not produce any more non-zero elements in the matrix; the factorization and solution can therefore use the same storage as the normal equations. Moreover, since the number of arithmetic operations grows only linearly with N , it is computationally feasible to solve the normal equations for any stretch of data.

¹¹ As indicated by the asterisk in the first derivative in Eq. (73), the differential in right ascension is a true arc, thus: $\partial \alpha_{*i} \equiv (\partial \alpha_i) \cos \delta_i$. The corresponding update in right ascension, i.e., the first element of the \mathbf{d}_i obtained by solving (57), is therefore $\Delta \alpha_i \cos \delta_i$.




Fig. 5. Structure of the attitude normal equations matrix N_{aa} for a cubic spline (order $M = 4$). The blocks of size 4×4 are indexed by (n, m) as in Eq. (33). The grey cells represent non-zero elements. Since the matrix is symmetric, the elements below the diagonal (in lighter grey) need not be stored.




Fig. 6. A natural break in the definition of the attitude spline occurs if there is a gap in the observations containing at least M knots, where M is the order of the spline. t_{last} is the time of the last observation before the gap, t_{first} the time of the first observation after the gap.

5.2.2. Segmentation of the data

Even though it is feasible to treat the complete set of normal equations for the attitude updating as a single system, it is desirable for several reasons to divide up the data temporally. This allows one to set up a very straightforward and efficient distributed attitude updating, simply by handing out the processing of different time segments to different computing nodes. Also the inspection of residuals in order to detect stretches of bad fit (caused, for example, by micrometeoroid impacts), and the subsequent reprocessing of these stretches, is greatly facilitated if it can be done on shorter data segments.

The spline model is capable of interpolating sensibly (if not accurately) over short data gaps. However, if the data gap contains at least M knots (with $M = 4$ for cubic splines), the two splines on each side of the gap become completely disconnected. This is illustrated in Fig. 6, where n and m are the left indices of, respectively, the last observation before the gap (t_{last}) and the first observation after the gap (t_{first}). $B_n(t)$ and $B_{m-M+1}(t)$ are the last and first B-splines whose coefficients depend on the observations before and after the gap. Clearly the two segments of the attitude spline are disconnected if $n < m - M + 1$ or $m - n \geq M$. We call this a natural attitude break.

In the absence of natural breaks, artificial ones can be introduced at suitable intervals by a simple method and without




Fig. 7. Illustrating the assignment of knots for the attitude update solutions in two consecutive segments K and $K+1$, with a break-point at knot index x and using $2L$ overlapping knot intervals.

any significant loss of accuracy. The idea is to make separate solutions for overlapping segments, as illustrated in Fig. 7. The segments use a common knot sequence $\{\tau_k\}$ that may extend over the whole length of the mission. Each segment K defines an attitude spline in the interval $[t_{\text{beg}}(K), t_{\text{end}}(K)]$, based on data with observation times in that same interval. The endpoints coincide with certain knots in such a way that $t_{\text{end}}(K) = \tau_{x+L}$ and $t_{\text{beg}}(K+1) = \tau_{x-L}$, where τ_x is the cross-over knot between segments K and $K+1$ and $2L$ the number of overlapping knot intervals. The anterior and posterior knots for each segment are also taken from the common knot sequence. The local character of the splines means that the resulting fit around τ_x is practically the same for the two segments, provided L is large enough. For cubic splines ($M = 4$) it is found that $L = 12$ is sufficient.

Each segment gives a system of normal equations (76) for the updates d_n to the attitude parameters a_n for a certain range of index n . For example, with reference to Fig. 7, in segment K updates are computed up to and including $n = x + L - 1$, while in segment $K+1$ updates are computed starting with $n = x - L - M + 1$. At least in the middle part of the overlap region, the updates for a given n should be essentially the same in the two segments. It therefore does not matter much which of the results is chosen. The mid-point is at index $n = x - M/2$ if M is even, or half-way between $x - (M+1)/2$ and $x - (M-1)/2$ if M is odd. We may therefore agree to use the solution from segment K to update a_n up to and including index $n = x - \lfloor M/2 \rfloor$, while the solution from segment $K+1$ is used starting with $n = x - \lfloor M/2 \rfloor + 1$. The important thing is that no n is missed by the updating, nor updated twice.

Once all the coefficients a_n have been updated from the different segmented solutions, the segmentation loses its meaning and can in principle be forgotten. For example, when evaluating the attitude at a specific time t , it does not matter to which segment that instant belonged. In the next iteration of the attitude update, a different segmentation can in principle be used.

The overlapping segments mean that a fraction of the observations need to be processed twice in the attitude updating. The fraction equals the ratio of the overlap length to the mean length of the segment, and increases the shorter the segments are made. For example, with a segment length of one day (it would not seem reasonable to have shorter segments) and a mean knot interval of 15 s, the fractional overlap for $L = 12$ is only 0.4%.

5.2.3. Calculation of partial derivatives

For the partial derivatives we obtain in analogy with Eq. (71)

$$-\frac{\partial R_l^{\text{AL}}}{\partial \mathbf{q}_l} = \mathbf{m}_l' \frac{\partial \mathbf{u}_l}{\partial \mathbf{q}_l} \sec \zeta_l, \quad -\frac{\partial R_l^{\text{AC}}}{\partial \mathbf{q}_l} = \mathbf{n}_l' \frac{\partial \mathbf{u}_l}{\partial \mathbf{q}_l}, \quad (77)$$

where the derivatives with respect to the attitude quaternion should be interpreted componentwise. Differentiating Eq. (11) and using that $d(C'u_l) = \mathbf{0}$, we have

$$\{d(S'u_l), 0\} = 2\{S'u_l, 0\}q_l^{-1}dq_l, \quad (78)$$

which after some manipulation gives

$$-\frac{\partial R_l^{\text{AL}}}{\partial q_l} = -2 \sec \zeta_l q_l \{S'n_l, 0\}, \quad -\frac{\partial R_l^{\text{AC}}}{\partial q_l} = 2q_l \{S'm_l, 0\}. \quad (79)$$

The derivatives with respect to the spline coefficients \mathbf{a}_n are obtained after multiplying the above expressions by $B_n(t_l)$, assuming that the normalization factor in Eq. (10) is close to unity (Sect. 5.2.4).

5.2.4. Constraining the attitude updating

Since the attitude is represented by a unit quaternion, its components should at all times satisfy $q_x^2 + q_y^2 + q_z^2 + q_w^2 = 1$. All four components are nevertheless needed, for while the magnitude of any of them can be inferred from the other three components, its sign cannot. The redundancy of the representation manifests itself in that the length of the quaternion cannot be determined from the observations. Indeed, as can be seen from Eqs. (11)–(12), applying an arbitrary non-zero scale factor to the attitude quaternion \mathbf{q} has no effect on the computed instrument angles, and is therefore unobservable. The attitude parameters \mathbf{a}_n are therefore also undefined with respect to a certain scale value. As a consequence, the normal matrix N_{aa} computed from Eq. (76) is singular, and constraints are needed for computing a unique solution.

The normalization in Eq. (10) was introduced to guarantee that the calculated quaternion is always of unit length, although, as we have seen, this is not strictly necessary for some of the subsequent calculations.¹²

Naively, one might expect that the coefficients \mathbf{a}_n could be scaled independently for each B-spline, i.e., that different scale factors could apply to each index n . This is not the case, however. At any time, the (non-normalized) attitude quaternion is a linear combination of M adjacent coefficients \mathbf{a}_n , and unless all four coefficients are scaled by exactly the same factor, the result will not be a simple scaling of the quaternion.¹³ Applying this argument to every observation time t_l , it is readily seen that the same scaling factor must be used for all the coefficients in any attitude segment without natural or artificial breaks. In principle, therefore, the attitude normal matrix for such a segment has a rank defect of 1 (that is, the rank is one less than the number of attitude parameters), and would only need a single constraint equation to become non-singular.

Numerical experiments, using Singular Value Decomposition (SVD; see, e.g., Golub & van Loan 1996) of the matrix N_{aa} computed from simulated observations over successively longer time intervals, indeed show the expected rank defect of 1 for intervals up to several hundred knots. That is, there is a clear gap (of several orders of magnitude) between the smallest singular value and the second smallest one. For longer time intervals

¹² On the other hand we have implicitly assumed $\|\mathbf{q}\| = 1$, for example in deriving Eq. (79).

¹³ Note that the normalization in Eq. (10) is effected by applying a normalization factor that is a continuous function of time. Since any given spline coefficient \mathbf{a}_n is used in a finite time interval (namely, the support of the corresponding B-spline), one cannot obtain a correct normalization for all t by scaling the spline coefficients by some value depending on n .

the gap gradually closes and the problem thus becomes ill-posed (Hansen 1998). Thus, any reasonably long time interval will in practice require some form of regularization rather than the application of just a single constraint.

The adopted solution method is a variant of the well known Tikhonov regularization (Hansen 1998). The objective function in Eq. (27) is modified to include a term depending on the deviation of the normalization factor in Eq. (10) from unity for each observation. We write the deviation as

$$D_l \equiv 1 - \left\| \sum_{n=\ell-M+1}^{\ell} \mathbf{a}_n B_n(t_l) \right\|^2 \quad (80)$$

and the modified objective function as

$$Q(s, \mathbf{a}, \mathbf{c}, \mathbf{g}) = \sum_l (R_l^2 + \lambda^2 D_l^2) W_l, \quad (81)$$

where λ is a small but non-zero regularization parameter. We have found that $\lambda = 10^{-3}$ to 10^{-2} gives a solution that is always numerically stable, and quite insensitive to the precise value of λ . As a result, the normal equations (76) become

$$\begin{aligned} \sum_{m=n-M+1}^{n+M-1} \left[\sum_{l \in L_n \cap L_m} \left(\frac{\partial R_l}{\partial \mathbf{a}_n} \frac{\partial R_l}{\partial \mathbf{a}_m'} + \lambda^2 \frac{\partial D_l}{\partial \mathbf{a}_n} \frac{\partial D_l}{\partial \mathbf{a}_m'} \right) W_l \right] \mathbf{d}_m \\ = - \sum_{l \in L_n} \left(\frac{\partial R_l}{\partial \mathbf{a}_n} R_l + \lambda^2 \frac{\partial D_l}{\partial \mathbf{a}_n} D_l \right) W_l. \end{aligned} \quad (82)$$

The required partial derivatives, obtained from Eq. (80), are

$$-\frac{\partial D_l}{\partial \mathbf{a}_n} = 2 \sum_{m=\ell-M+1}^{\ell} \mathbf{a}_m B_m(t_l) B_n(t_l) \simeq 2 \mathbf{q}(t_l) B_n(t_l), \quad (83)$$

where the approximation makes use of the fact that the normalization factor in Eq. (10) is close to unity.

5.2.5. Estimating the excess attitude noise

The excess attitude noise $\epsilon_a(t)$ introduced in Eq. (28) accounts for modelling errors in the attitude representation. Such errors could be caused for example by (unmodelled) micrometeoroid impacts, ‘clanks’ due to sudden redistributions of satellite inertia, propellant sloshing, thruster noise, or mechanical vibrations (Appendix D.4). Due to the cubic spline representation, any localized effect that cannot be fitted by the spline will result in systematic residuals that span over a few consecutive knot intervals. Indeed, discontinuities in the rate (e.g., from micrometeoroid impacts) or angle (e.g., from clanks) produce characteristic patterns of residuals that can be used to identify such events. A significant effort will be devoted to the possibly semi-manual and interactive process of finding these events. When identified, they can be handled for example by modifying the knot sequence (Sect. 5.2.6). But even after this process, the model will be imperfect due to for example high-frequency thruster noise.¹⁴ Similarly, there will be a large number of impacts that are too small to be individually recognized; collectively they add some unmodelled attitude errors, which $\epsilon_a(t)$ may account for. However, it should be noted that $\epsilon_a(t)$ does *not* include any

¹⁴ ‘High-frequency’ here means roughly the range $1/2\Delta\tau \simeq 0.03$ Hz to 0.2 Hz, where $\Delta\tau$ is the typical spline knot interval (~ 15 s); lower frequencies are absorbed by the spline and higher frequencies are smoothed out by the integration across the CCD.

component of the observation noise (principally from CCD photon noise), nor is it an estimation of the attitude uncertainty (cf. Sect. 6.3).

Three components of the excess noise, designated $\epsilon_{AL}(t)$, $\epsilon_{ACP}(t)$, and $\epsilon_{ACF}(t)$, need to be derived independently of each other, representing modelling errors in the AL attitude, the AC attitude of the preceding field of view, and the AC attitude in the following field of view. The three components are statistically nearly independent thanks to the way the attitude measurements are made, and the fact that the basic angle is not far from 90° .

The algorithm to estimate $\epsilon_a(t)$ (for $a = AL, ACP$ or ACF) is based on a simple statistical processing of the residuals R_l derived in the source updating. The time line t is divided into ‘buckets’ $[t_j, t_{j+1})$ such that each bucket (j) will contain a sufficient number of observations, also in the AC direction. The size (duration) of a bucket should be several knot intervals for the attitude spline, but the boundaries t_j need not in any other way be related to the attitude knot sequence. One set of buckets is needed for each attitude component (AL, ACP, ACF). Let $l \in ja$ signify that observation l belongs to bucket j and attitude component a . After having completed the source updating for all primary sources, the excess attitude noise in bucket j is estimated as

$$\epsilon_a^2(t_j \leq t < t_{j+1}) = \max \left(0, F_{0.68}(R_l^2 - \sigma_l^2 - \epsilon_i^2) \right), \quad (84)$$

where $F_{0.68}()$ is the 68% quantile (68th percentile) of the argument values. It is important to note that the downweighting factors w_l determined during the source updating are not used here to eliminate possible outliers; this function is instead taken care of by using the quantile to compute a robust estimate of the typical excess variance in the attitude bucket. This means that if the ‘outliers’ detected by the source update were actually caused by a stretch of bad attitude, then this will be recognized by a large value of the quantile in Eq. (84), and consequently by an increased ϵ_a^2 .

In the subsequent attitude update, the downweighting factors w_l are re-computed based on the residual from the previous source update but with a value for the total variance, $\sigma_l^2 + \sigma_i^2 + \epsilon_a^2$, using the newly estimated ϵ_a^2 . Thus, only the ‘true’ outliers – that are not due to the bad attitude – are now downweighted. The data may thus contribute to the attitude updating even if they had been flagged as outliers in the preceding source updating.

The functions $\epsilon_a(t)$ are obviously an extremely useful diagnostic for the progress of the AGIS iterations as well as (after convergence) for the quality of the attitude modelling and data. They can be plotted as a function of time, and the quantity of data is such that human inspection is feasible. They are also needed for setting the detection threshold for micrometeoroid impacts.

The accumulation of statistics in the buckets is best done in parallel with the source updating, when the residuals are readily at hand. One remaining problem is how to compute the quantile in Eq. (84) in an efficient way, without having to store billions of residuals. Indeed, exact calculation of quantiles would require to store all the values $R_l^2 - \sigma_l^2 - \epsilon_i^2$ in a bucket before the quantile can be computed. However, if we are content with an approximate estimate of the quantile, there are a number of sequential estimation algorithms available that only need to store a much smaller amount of data per bucket, see for example Greenwald & Khanna (2001), Gilbert et al. (2002) and references therein. We have chosen to use the Incremental Quantile estimation algorithm due to Chambers et al. (2006).

5.2.6. Initialization of the attitude parameters

An approximate estimate of the attitude is already provided by the initial data treatment (IDT) preceding the astrometric solution. This may be given as a discrete time series, for example one quaternion every second of time. The first time the attitude update is executed for a certain time interval, a regular knot sequence is set up and the B-spline coefficients \mathbf{a}_n in Eq. (10) are determined by a least-squares fit. For a given time series of attitude estimates, this is a linear problem and therefore easily solved. The resulting initial attitude $\mathbf{a}^{(0)}$ is used in the first source update (S) and subsequently improved by the attitude update process (A) as part of the AGIS iteration scheme.

By default, a regular knot sequence is adopted, i.e., the knot interval $\Delta\tau = \tau_{n+1} - \tau_n$ is taken to be more or less constant. Given the endpoints $t_{\text{beg}}, t_{\text{end}}$ of a data segment, the knots are set up at regular intervals respecting a given maximum value of $\Delta\tau$ (of the order of 5 to 20 s). The assignment of knots must also take into account the need for anterior and posterior knots, as discussed in Appendix B, and in the case of segmented data, the overlapping knots as discussed in Sect. 5.2.2.

Occasionally the knot sequence needs to be redefined as a result of the adjustment process. Possible causes could be:

- If the spline is not flexible enough to accurately model the data, it may be necessary to decrease the maximum allowed $\Delta\tau$.
- Conversely, overfitting of the data may require the maximum allowed $\Delta\tau$ to be increased.
- Locally, a scarcity of accurate data or a short gap could make it necessary to remove some knots or introduce a natural break in the attitude representation (Sect. 5.2.2).
- Very locally, a bad fit may result from a micrometeoroid hit causing an almost instantaneous change in the angular velocity of the satellite. This may be dealt with by introducing multiple knots at the appropriate instants (Appendix D.4 and B.3).

Having redefined the knot sequence, it is necessary to re-initialize the spline coefficients \mathbf{a}_n , which must now refer to the new knot sequence. This is most easily done by evaluating $\mathbf{q}(t)$ for a regular time series, with a sampling interval much smaller than $\Delta\tau$ (e.g., 1 s), and fitting the new spline to the time series.

5.3. Calibration updating (C)

The geometric instrument model is given in Sect. 3.4. We assume here the generic calibration model in Eqs. (20)–(21), in which the parameters are indexed by rs . In the calibration update block (C) the calibration parameters \mathbf{c} are improved by solving the third sub-system in Eq. (38), i.e., the normal equations

$$\left[\sum_l \frac{\partial R_l}{\partial \mathbf{c}} \frac{\partial R_l}{\partial \mathbf{c}'} W_l \right] \mathbf{d}_c = - \sum_l \frac{\partial R_l}{\partial \mathbf{c}} R_l(s, \mathbf{a}, \mathbf{c}, \mathbf{g}) W_l. \quad (85)$$

The residuals in the right-hand side are computed from the parameters values in the current or preceding iteration according to the discussion in Sect. 4.6. Because the calibration model is linear, the partial derivatives are uniquely determined by the observation index l ,

$$\frac{\partial R_l}{\partial c_{rs}} = \frac{\partial \eta_l^{\text{obs}}}{\partial c_{rs}} = \begin{cases} \Phi_{rs}(l) & \text{if } l \in rs, \\ 0 & \text{otherwise.} \end{cases} \quad (86)$$

In the normal matrix, the element with subscripts $(rs)_1$ and $(rs)_2$ is non-zero only if there is at least one observation l such that

$l \in (rs)_1$ and $l \in (rs)_2$. Depending on how the calibration parameters are grouped into sets with no common observations (for example according to the CCD/gate combination; cf. Sect. 4.3), the normal matrix will therefore be block-diagonal, which the calibration updating takes advantage of in order to save computations. It also facilitates distributed processing.

Since the weights W_l are fixed from the preceding source and attitude updating processes, the update \mathbf{d}_c can be calculated in a single direct solution, using the robust Cholesky decomposition (Appendix C). However, due to the degeneracy between for example the large-scale and small-scale AL calibration parameters, this will produce an arbitrary feasible solution $\tilde{\mathbf{d}}_c$, which does not necessarily satisfy the constraints in Eq. (22). The constrained update is obtained as

$$\mathbf{d}_c = \tilde{\mathbf{d}}_c - \mathbf{C}(\mathbf{C}'\mathbf{C})^{-1}\mathbf{C}'\tilde{\mathbf{d}}_c, \quad (87)$$

whereupon the updated \mathbf{c} can be computed.

The above-mentioned degeneracy among the calibration parameters means that the normal matrix calculated according to Eq. (85) is singular, which seems to contradict our assumption (Sect. 4.3) that N_{cc} is positive-definite. However, if \mathbf{C} spans the null space of N_{cc} , as it should for a properly formulated set of constraints, then it can be seen that Eq. (87) gives the same result as solving the updates with the modified normal matrix $N_{cc} + \lambda^2\mathbf{C}\mathbf{C}'$, which is positive definite for any $\lambda \neq 0$. Thus, the procedure outlined above is equivalent to solving the constrained system with positive-definite matrix.

5.4. Global updating

An arbitrary number of global parameters may be solved for in the AGIS system. Global parameters should be defined in such a way that their default values, equal to zero, correspond to the baseline solution. By not solving for the globals, we implicitly set them to zero, resulting in the baseline solution. For example, we have a very high confidence in General Relativity, which in the parametrized post-Newtonian (PPN) formalism implies the parameter $\gamma = 1$. The global parameter related to the gravitational deflection of light should therefore not be γ itself, but for example the parameter g_0 in

$$\gamma = 1 + g_0. \quad (88)$$

That is, $g_0 = 0$ corresponds to the baseline case of General Relativity. The global parameter vector is $\mathbf{g} = (g_0, g_1, \dots)'$.

The normal equations for the update \mathbf{d}_g to the global parameter vector are

$$\left[\sum_l \frac{\partial R_l}{\partial \mathbf{g}} \frac{\partial R_l}{\partial \mathbf{g}'} W_l \right] \mathbf{d}_g = - \sum_l \frac{\partial R_l}{\partial \mathbf{g}} R_l(s, \mathbf{a}, \mathbf{c}, \mathbf{g}) W_l, \quad (89)$$

where the sums are taken over all the observations l , and the statistical weights W_l follow from the preceding source and attitude updates. The partial derivatives in Eq. (89) are computed in exact analogy with Eqs. (71)–(72) for the source updating, only with \mathbf{g} replacing \mathbf{s}_i . The calculation of $\partial \mathbf{u}_l / \partial \mathbf{g}'$ is not detailed here.

In the simple iteration scheme (Sect. 4.5), the inclusion of g_0 representing the PPN parameter γ considerably slows down the convergence of the astrometric solution. As explained by Hobbs et al. (2010), this behaviour is caused by the relatively strong correlation between the gravitational light deflection by the Sun (proportional to $1 + \gamma$, and directed away from the Sun) and trigonometric parallax (directed towards the solar-system barycentre, never far from the Sun). Hobbs et al. (2010) found

that the convergence rate could be restored by the introduction of a pseudo-parameter g_1 representing a global shift of all parallaxes. (The update to this parameter is solved in each iteration but never applied – its value remains at zero and the converged values of all the other parameters are unaffected; hence the prefix ‘pseudo’.) It was later found that this artefact is not needed when using the conjugate gradients scheme, which gives roughly the same rate of convergence whether or not g_1 is included.

6. Auxiliary processes

In this section we describe some auxiliary processes that are not necessarily part of the astrometric solution as such, but nevertheless needed in order to construct the astrometric catalogue. They concern the definition of the reference system for the source positions and proper motions by means of the frame rotator (Sect. 6.1), the selection of primary sources (Sect. 6.2), and the computation of the standard uncertainties and correlations of the astrometric parameters (Sect. 6.3).

6.1. Frame rotator

As explained in Sect. 4.4, the measurement principle of Gaia results in a system of positions and proper motions that is essentially undefined with respect to an arbitrary (small) offset in the orientation and spin of the reference frame. As a consequence, the normal matrix \mathbf{N} is in principle singular with a rank defect of 6.

While the solution of rank-deficient problems in general requires special attention to the singularities, for example by adding constraints to avoid numerical instability, no such complication arises here because of the way AGIS works. Basically, a solution is found by iterating between the source and attitude updates (the calibration and global updates play no role here because they are to first order independent of the reference frame). When the sources are updated, the reference frame is in reality set by the (then assumed) attitude; similarly, when the attitude is updated, the frame is set by the (then assumed) source parameters. In terms of the matrix formulation of Sect. 4.5 this is equivalent to the statement that the preconditioner \mathbf{K} is non-singular. The end result is that AGIS converges to a solution with both the source and attitude parameters expressed in the same, but largely arbitrary, reference frame.

The intention is however that the final source parameters (positions and proper motions) shall be expressed in a celestial reference frame that represents, as closely as possible, the International Celestial Reference System (ICRS). For consistency, it is moreover necessary that the attitude parameters are expressed in exactly the same frame as the source parameters. It is the task of the frame rotator to accomplish this. A similar process was used to align the Hipparcos Catalogue with the extragalactic reference frame (Lindegren & Kovalevsky 1995).

In the following we start with the rigorous definition of the rotation correction, then derive a linear approximation applicable to the small corrections that we have in practice. Finally, we discuss the determination of the rotation parameters and their application in the AGIS iteration scheme.

6.1.1. Relation between the ICRS and AGIS frames

Ideally, the astrometric solution should result in parameters that are expressed in the BCRS (Sect. 3.1), whose axes are aligned with the ICRS here represented by $\mathbf{C} = [\mathbf{X} \ \mathbf{Y} \ \mathbf{Z}]$. However, due

to the in principle undefined reference frame of AGIS, the astrometric solution is in effect expressed relative to a slightly different triad, which we denote $\tilde{\mathbf{C}} = [\tilde{\mathbf{X}} \tilde{\mathbf{Y}} \tilde{\mathbf{Z}}]$. The two reference systems, which for simplicity will be referred to as the ICRS and AGIS frames, are related by a time-dependent spatial rotation given by the quaternion $\mathbf{f}(t)$; thus the coordinates of the arbitrary (fixed) vector \mathbf{v} are transformed according to the frame rotation formula

$$\{\mathbf{C}'\mathbf{v}, 0\} = \mathbf{f}(t)^{-1} \{\tilde{\mathbf{C}}'\mathbf{v}, 0\} \mathbf{f}(t) \quad (90)$$

(cf. Eq. A.15). Due to the kinematical constraints of the AGIS solution, $\mathbf{f}(t)$ describes a uniform spin motion of the two frames with respect to each other.

For consistency with Lindegren & Kovalevsky (1995) we parametrize $\mathbf{f}(t)$ by means of two vectors $\boldsymbol{\varepsilon}$ and $\boldsymbol{\omega}$ representing *corrections* to the orientation and spin of the AGIS frame. More precisely, the parameters of the frame rotator are the six coordinates of the vectors in the AGIS frame at some adopted frame rotator epoch t_{fr} (not necessarily the same as the reference epoch t_{ep} of the astrometric parameters). These coordinates are denoted $\varepsilon_{\tilde{x}}, \varepsilon_{\tilde{y}}, \varepsilon_{\tilde{z}}, \omega_{\tilde{x}}, \omega_{\tilde{y}}, \omega_{\tilde{z}}$; according to our kinematical assumption they are strictly constant numbers. For the arbitrary epoch t the frame rotator quaternion is, therefore,

$$\mathbf{f}(t) = \mathbf{Q}[(t - t_{\text{fr}})\tilde{\mathbf{C}}'\boldsymbol{\omega}]\mathbf{Q}(\tilde{\mathbf{C}}'\boldsymbol{\varepsilon}), \quad (91)$$

where \mathbf{Q} is the function introduced by Eq. (A.12). Equations (90) and (91) provide the basis for the rigorous transformation of any data between the two frames, given the rotation parameters $\tilde{\mathbf{C}}'\boldsymbol{\varepsilon} = [\varepsilon_{\tilde{x}} \varepsilon_{\tilde{y}} \varepsilon_{\tilde{z}}]'$ and $\tilde{\mathbf{C}}'\boldsymbol{\omega} = [\omega_{\tilde{x}} \omega_{\tilde{y}} \omega_{\tilde{z}}]'$.

While the above expressions are strictly valid for arbitrarily large rotation parameters, we have in practice $\|\boldsymbol{\varepsilon}\|, \|(t - t_{\text{fr}})\boldsymbol{\omega}\| < 20 \text{ mas} \approx 10^{-7} \text{ rad}$, at least in the final iterations of AGIS. This means that second-order terms are completely negligible ($< 0.002 \mu\text{as}$). To first order we have

$$\mathbf{f}(t) \approx \mathbf{Q}[\tilde{\mathbf{C}}'\boldsymbol{\varepsilon} + (t - t_{\text{fr}})\tilde{\mathbf{C}}'\boldsymbol{\omega}], \quad (92)$$

and the vector part of Eq. (90) becomes, to the same approximation,

$$\begin{bmatrix} v_X \\ v_Y \\ v_Z \end{bmatrix} = \begin{bmatrix} v_{\tilde{X}} \\ v_{\tilde{Y}} \\ v_{\tilde{Z}} \end{bmatrix} + \begin{bmatrix} 0 & -v_{\tilde{Z}} & +v_{\tilde{Y}} \\ +v_{\tilde{Z}} & 0 & -v_{\tilde{X}} \\ -v_{\tilde{Y}} & +v_{\tilde{X}} & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_{\tilde{X}} + (t - t_{\text{fr}})\omega_{\tilde{X}} \\ \varepsilon_{\tilde{Y}} + (t - t_{\text{fr}})\omega_{\tilde{Y}} \\ \varepsilon_{\tilde{Z}} + (t - t_{\text{fr}})\omega_{\tilde{Z}} \end{bmatrix}. \quad (93)$$

6.1.2. Transformation of the astrometric parameters

Let $\tilde{\alpha}, \tilde{\delta}, \tilde{\mu}_{\alpha*}, \tilde{\mu}_{\delta}$ be the position and proper motion parameters for a source as derived in AGIS, that is referring to $\tilde{\mathbf{C}}$. For brevity we omit here the source index (i), and do not consider the parallax ϖ_i and radial proper motion $\mu_{r,i}$ which are independent of the frame orientation. In analogy with Eq. (5) we have the normal triad $[\tilde{\mathbf{p}} \tilde{\mathbf{q}} \mathbf{r}]$ with respect to the AGIS frame, where \mathbf{r} is the barycentric direction to the source at time t_{ep} , $\tilde{\mathbf{p}} = \langle \tilde{\mathbf{X}} \times \mathbf{r} \rangle$ and $\tilde{\mathbf{q}} = \mathbf{r} \times \tilde{\mathbf{p}}$; its coordinates in the AGIS frame are given by the columns of the matrix

$$\tilde{\mathbf{C}}'[\tilde{\mathbf{p}} \tilde{\mathbf{q}} \mathbf{r}] = \begin{bmatrix} -\sin \tilde{\alpha} & -\sin \tilde{\delta} \cos \tilde{\alpha} & \cos \tilde{\delta} \cos \tilde{\alpha} \\ \cos \tilde{\alpha} & -\sin \tilde{\delta} \sin \tilde{\alpha} & \cos \tilde{\delta} \sin \tilde{\alpha} \\ 0 & \cos \tilde{\delta} & \sin \tilde{\delta} \end{bmatrix}. \quad (94)$$

At the source reference epoch t_{ep} the direction cosines of \mathbf{r} are related by the frame rotation in Eq. (91); thus

$$\{\mathbf{C}'\mathbf{r}, 0\} = \mathbf{f}(t_{\text{ep}})^{-1} \{\tilde{\mathbf{C}}'\mathbf{r}, 0\} \mathbf{f}(t_{\text{ep}}). \quad (95)$$

From $\mathbf{C}'\mathbf{r} = [r_x \ r_y \ r_z]'$ the position parameters in the ICRS frame are obtained as

$$\alpha = \text{atan2}(r_y, r_x), \quad \delta = \text{atan2}\left(r_z, \sqrt{r_x^2 + r_y^2}\right). \quad (96)$$

The transformation of the proper motion components is a bit more complicated, as they are expressed with respect to axes that are physically (slightly) different in the two frames, viz., $\tilde{\mathbf{p}}, \tilde{\mathbf{q}}$ in the AGIS frame, and \mathbf{p}, \mathbf{q} in the ICRS frame. However, the time derivative (at epoch t_{ep}) of the barycentric direction to the source is a fixed vector in space, known as the proper motion vector. In a kinematically non-rotating system it can be written

$$\boldsymbol{\mu} = \mathbf{p}\mu_{\alpha*} + \mathbf{q}\mu_{\delta} = \tilde{\mathbf{p}}\tilde{\mu}_{\alpha*} + \tilde{\mathbf{q}}\tilde{\mu}_{\delta} - \boldsymbol{\omega} \times \mathbf{r}, \quad (97)$$

where the last term is the correction for the spin of the AGIS frame. The coordinates of the proper motion vector in the two frames,

$$\mathbf{C}'\boldsymbol{\mu} = \mathbf{C}'\mathbf{p}\mu_{\alpha*} + \mathbf{C}'\mathbf{q}\mu_{\delta} \quad (98)$$

and

$$\tilde{\mathbf{C}}'\boldsymbol{\mu} = \tilde{\mathbf{C}}'\tilde{\mathbf{p}}\tilde{\mu}_{\alpha*} + \tilde{\mathbf{C}}'\tilde{\mathbf{q}}\tilde{\mu}_{\delta} - (\tilde{\mathbf{C}}'\boldsymbol{\omega}) \times (\tilde{\mathbf{C}}'\mathbf{r}) \quad (99)$$

are related by a frame rotation analogous to Eq. (95),

$$\{\mathbf{C}'\boldsymbol{\mu}, 0\} = \mathbf{f}(t_{\text{ep}})^{-1} \{\tilde{\mathbf{C}}'\boldsymbol{\mu}, 0\} \mathbf{f}(t_{\text{ep}}). \quad (100)$$

From Eq. (97) the proper motion components in the ICRS frame are then

$$\mu_{\alpha*} = \mathbf{p}'\boldsymbol{\mu} = (\mathbf{C}'\mathbf{p})'(\mathbf{C}'\boldsymbol{\mu}), \quad \mu_{\delta} = \mathbf{q}'\boldsymbol{\mu} = (\mathbf{C}'\mathbf{q})'(\mathbf{C}'\boldsymbol{\mu}). \quad (101)$$

For given $(\tilde{\alpha}, \tilde{\delta}, \tilde{\mu}_{\alpha*}, \tilde{\mu}_{\delta})$ and $(\tilde{\mathbf{C}}'\boldsymbol{\varepsilon}, \tilde{\mathbf{C}}'\boldsymbol{\omega})$, the sequence of calculations is therefore:

1. Calculate $\mathbf{f}(t_{\text{ep}})$ by Eq. (91);
2. Calculate $\tilde{\mathbf{C}}'[\tilde{\mathbf{p}} \tilde{\mathbf{q}} \mathbf{r}]$ by Eq. (94);
3. Calculate $\mathbf{C}'\mathbf{r}$ by Eq. (95) and $\mathbf{C}'\boldsymbol{\mu}$ by Eq. (100);
4. Calculate α and δ by Eq. (96) and $\mathbf{C}'\mathbf{p}$ and $\mathbf{C}'\mathbf{q}$ by Eq. (5);
5. Calculate $\mu_{\alpha*}$ and μ_{δ} by Eq. (101).

As we have not employed the approximations in Eqs. (92)–(93), these transformations are rigorous.

6.1.3. Transformation of the attitude parameters

In analogy with Eq. (9) the attitude quaternion $\tilde{\mathbf{q}}(t)$ derived in AGIS defines the transformation from the AGIS frame $\tilde{\mathbf{C}}$ to the SRS \mathbf{S} as a function of time; thus for the arbitrary vector \mathbf{v}

$$\{\mathbf{S}'\mathbf{v}, 0\} = \tilde{\mathbf{q}}(t)^{-1} \{\tilde{\mathbf{C}}'\mathbf{v}, 0\} \tilde{\mathbf{q}}(t). \quad (102)$$

Solving $\{\tilde{\mathbf{C}}'\mathbf{v}, 0\}$ and inserting into Eq. (90) yields

$$\{\mathbf{C}'\mathbf{v}, 0\} = \mathbf{f}(t)^{-1} \tilde{\mathbf{q}}(t) \{\mathbf{S}'\mathbf{v}, 0\} \tilde{\mathbf{q}}(t)^{-1} \mathbf{f}(t). \quad (103)$$

Comparison with Eq. (9) shows that the corrected attitude is given by

$$\mathbf{q}(t) = \mathbf{f}(t)^{-1} \tilde{\mathbf{q}}(t). \quad (104)$$

In practice the AGIS attitude $\tilde{\mathbf{q}}(t)$ is expressed in terms of B-splines by means of coefficients $\tilde{\mathbf{a}}_n$ as in Eq. (10). The result of the time-dependent transformation by $\mathbf{f}(t)^{-1}$ in Eq. (104) cannot, in general, be exactly represented by means of B-splines. However, since the transformation is changing extremely slowly in comparison with the duration of the support of each B-spline ($\sim 1 \text{ min}$), and also the changes of \mathbf{a} from knot to knot are very

small, we make a negligible error by transforming the coefficients instead of the attitude quaternion. Thus we use

$$\mathbf{a}_n = \mathbf{f}(\bar{t}_n)^{-1} \tilde{\mathbf{a}}_n, \quad (105)$$

where

$$\bar{t}_n = \frac{1}{2}(\tau_n + \tau_{n+M}) \quad (106)$$

is the time half-way through the support of $B_n(t)$.

6.1.4. Determination of the frame rotator parameters

The parameters $\tilde{\mathbf{C}}'\boldsymbol{\varepsilon} = [\varepsilon_{\tilde{X}}, \varepsilon_{\tilde{Y}}, \varepsilon_{\tilde{Z}}]'$ and $\tilde{\mathbf{C}}'\boldsymbol{\omega} = [\omega_{\tilde{X}}, \omega_{\tilde{Y}}, \omega_{\tilde{Z}}]'$ are determined by a weighted least-squares solution, using as input the differences in positions and proper motions, for a subset of the sources, between the AGIS results and a priori data. Three kinds of sources may be used for this purpose:

- A subset S_{NR} of the primary sources can be assumed to define a kinematically non-rotating celestial frame. Typically this subset will contain some 10^5 to 10^6 quasars and point-like galactic nuclei, mainly identified from ground-based surveys and photometric criteria. This subset effectively determines $\boldsymbol{\omega}$.
- A subset S_{P} of S_{NR} in addition have positions $(\hat{\alpha}, \hat{\delta})$ accurately determined with respect to the ICRS by means that are completely independent of Gaia. Typically it will contain the optical counterparts of extragalactic objects with accurate positions from radio interferometry (VLBI). Due to the cosmological acceleration effect described below it is necessary to assign an epoch t_{P} to each such position. This subset effectively determines $\boldsymbol{\varepsilon}$.
- The third subset S_{PM} consists of primary sources that do not a priori belong to the non-rotating subset, but have positions and/or proper motions that are accurately determined with respect to the ICRS independent of Gaia. This could include radio stars observed by VLBI, or stars whose absolute proper motions have been determined by some other means. The astrometric parameters of a source in this subset are denoted $\hat{\alpha}, \hat{\delta}, \hat{\mu}_{\alpha*}, \hat{\mu}_{\delta}$ and refer to the epoch t_{PM} (the parallax is irrelevant here, as it is identical in both frames). It is not expected that this subset will contribute very significantly to the determination of $\boldsymbol{\varepsilon}$ and/or $\boldsymbol{\omega}$, but they are included in the discussion below since they may provide important consistency checks.

In the following we derive the appropriate observation equations for the different kinds of sources. The derivation assumes that the frame rotator parameters are numerically small so that the linear approximation in Eq. (93) applies. For the (weighted) least-squares estimation of the frame rotator parameters it is, furthermore, necessary to assign the appropriate statistical weights to the observations and to have procedures for identifying and handling outliers. These issues are however not discussed here.

In the frame rotator solution, the six parameters must be complemented by three more parameters a_X, a_Y, a_Z taking into account the acceleration of the solar-system barycentre in a cosmological frame (Bastian 1995; Gwinn et al. 1997; Kopeikin & Makarov 2006). Such an acceleration, by the vector $\boldsymbol{\alpha}$, will cause a systematic ‘streaming’ (dipole) pattern of the apparent proper motions of extragalactic objects, described by

$$\boldsymbol{\mu}_0 = (\mathbf{I} - \mathbf{r}\mathbf{r}')\boldsymbol{\alpha}. \quad (107)$$

Here \mathbf{r} is the direction to the source and $\boldsymbol{\alpha} = \boldsymbol{\alpha}/c$ to first order in c^{-1} , where c is the speed of light. The galactocentric acceleration of the solar-system barycentre by $\|\boldsymbol{\alpha}\| \approx 2 \times 10^{-10} \text{ m s}^{-2}$ is

expected to produce a proper motion pattern with an amplitude of $\|\boldsymbol{\alpha}\| \approx 4 \mu\text{as yr}^{-1}$, which Gaia should be able to detect given a sufficient number of quasars among the primary sources.¹⁵ The additional parameters introduced in the frame rotator solution are the components of $\boldsymbol{\alpha}$ in the ICRS, or $[a_X a_Y a_Z] = \mathbf{C}'\boldsymbol{\alpha}$; they may be expressed in the same unit as the proper motions.

Observation equations for a source in S_{NR} . A kinematically non-rotating source should only have an apparent proper motion due to the cosmological acceleration. Equating $\boldsymbol{\mu}$ in (97) with $\boldsymbol{\mu}_0$ from Eq. (107) and taking the scalar products with $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$ results in the two observation equations

$$\left. \begin{aligned} \tilde{\mathbf{p}}'\boldsymbol{\alpha} + \tilde{\mathbf{q}}'\boldsymbol{\omega} &= \tilde{\mu}_{\alpha*} \\ \tilde{\mathbf{q}}'\boldsymbol{\alpha} - \tilde{\mathbf{p}}'\boldsymbol{\omega} &= \tilde{\mu}_{\delta} \end{aligned} \right\}, \quad (108)$$

where we have used the scalar triple product rule¹⁶ for the terms including $\boldsymbol{\omega}$. These equations are linear in the unknown acceleration and spin parameters, and the coefficients are the known coordinates of $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{q}}$ in either $\tilde{\mathbf{C}}$ or \mathbf{C} (to the adopted approximation the coefficients are the same in the two frames).

Observation equations for a source in S_{P} . In order to compare positions it is necessary to choose an epoch at which to make the comparison. At the chosen epoch of comparison, t , the barycentric coordinate direction of the source is, to first order in the proper motion,

$$\begin{aligned} \bar{\mathbf{u}}_{\text{B}}(t) &= \mathbf{r}(t_{\text{P}}) + (t - t_{\text{P}})\boldsymbol{\mu}_0 \\ &= \mathbf{r}(t_{\text{ep}}) + (t - t_{\text{ep}})(\tilde{\mathbf{p}}'\tilde{\boldsymbol{\mu}}_{\alpha*} + \tilde{\mathbf{q}}'\tilde{\boldsymbol{\mu}}_{\delta} - \boldsymbol{\omega} \times \mathbf{r}). \end{aligned} \quad (109)$$

In the first equality we have made the assumption that the source has the apparent proper motion $\boldsymbol{\mu}_0$ when observed in the ICRS frame; the second uses the same proper motion vector derived from the AGIS data according to Eq. (97). If we now compute the coordinates of $\bar{\mathbf{u}}_{\text{B}}(t)$ in \mathbf{C} (using the first equality) and $\tilde{\mathbf{C}}$ (using the second equality), they must be related according to Eq. (93). Resolving the coordinate differences along α and δ we obtain the observation equations

$$\left. \begin{aligned} (t - t_{\text{P}})\tilde{\mathbf{p}}'\boldsymbol{\alpha} + \tilde{\mathbf{q}}'(\boldsymbol{\varepsilon} + (t - t_{\text{fr}})\boldsymbol{\omega}) &= \Delta\alpha^* \\ (t - t_{\text{P}})\tilde{\mathbf{q}}'\boldsymbol{\alpha} - \tilde{\mathbf{p}}'(\boldsymbol{\varepsilon} + (t - t_{\text{fr}})\boldsymbol{\omega}) &= \Delta\delta \end{aligned} \right\}, \quad (110)$$

where

$$\begin{bmatrix} \Delta\alpha^* \\ \Delta\delta \end{bmatrix} = \begin{bmatrix} \tilde{p}_{\tilde{X}} & \tilde{p}_{\tilde{Y}} & \tilde{p}_{\tilde{Z}} \\ \tilde{q}_{\tilde{X}} & \tilde{q}_{\tilde{Y}} & \tilde{q}_{\tilde{Z}} \end{bmatrix} \begin{bmatrix} \cos \tilde{\delta} \cos \tilde{\alpha} - \cos \hat{\delta} \cos \hat{\alpha} \\ \cos \tilde{\delta} \sin \tilde{\alpha} - \cos \hat{\delta} \sin \hat{\alpha} \\ \sin \tilde{\delta} - \sin \hat{\delta} \end{bmatrix} + (t - t_{\text{ep}}) \begin{bmatrix} \tilde{\mu}_{\alpha*} \\ \tilde{\mu}_{\delta} \end{bmatrix}. \quad (111)$$

The observation equations in proper motion are of course the same as in Eq. (108).

Returning to the choice of comparison epoch t , it is clear that the result in terms of the estimated frame rotator parameters should in principle not depend on this choice. However, that will only be the case if the statistical correlations among the data are

¹⁵ The expected acceleration due to, for example, the Andromeda galaxy or the Shapley Supercluster (Proust et al. 2006) are at most $\sim 10^{-3}$ of the galactocentric acceleration. On the other hand, nearby massive galactic objects and large-scale deviations of the galactic potential from axisymmetry could conceivably produce a larger deviation in the direction of the total acceleration.

¹⁶ $\mathbf{a}'(\mathbf{b} \times \mathbf{c}) = \mathbf{b}'(\mathbf{c} \times \mathbf{a}) = \mathbf{c}'(\mathbf{a} \times \mathbf{b})$

taken into account in the least-squares estimation. Otherwise one should choose t to minimize these correlations. From Eq. (111) it is seen that the right-hand sides of Eq. (108) and (110) are uncorrelated if $t = t_{\text{ep}}$, provided that the position and proper motion errors in AGIS are also uncorrelated (which is generally the case if t_{ep} is appropriately chosen, i.e., close to mid-mission). Consequently we suggest using $t = t_{\text{ep}}$.

Observation equations for a source in S_{PM} . Here it will be necessary to distinguish two cases depending on how the proper motions have been measured. If $\hat{\mu}_{\alpha^*}$, $\hat{\mu}_{\delta}$ are the proper motion components of a source measured relative to background quasars, then the observation equations are simply

$$\left. \begin{aligned} \tilde{\mathbf{p}}' \mathbf{a} + \tilde{\mathbf{q}}' \boldsymbol{\omega} &= \tilde{\mu}_{\alpha^*} - \hat{\mu}_{\alpha^*} \\ \tilde{\mathbf{q}}' \mathbf{a} - \tilde{\mathbf{p}}' \boldsymbol{\omega} &= \tilde{\mu}_{\delta} - \hat{\mu}_{\delta} \end{aligned} \right\}. \quad (112)$$

As expected, the equations simplify to Eq. (108) in case the measured proper motion is zero. The observation equations obtained from the comparison of positions are the same as in Eq. (110), but with right-hand side

$$\begin{aligned} \begin{bmatrix} \Delta\alpha^* \\ \Delta\delta \end{bmatrix} &= \begin{bmatrix} \tilde{p}_X & \tilde{p}_Y & \tilde{p}_Z \\ \tilde{q}_X & \tilde{q}_Y & \tilde{q}_Z \end{bmatrix} \begin{bmatrix} \cos \tilde{\delta} \cos \tilde{\alpha} - \cos \hat{\delta} \cos \hat{\alpha} \\ \cos \tilde{\delta} \sin \tilde{\alpha} - \cos \hat{\delta} \sin \hat{\alpha} \\ \sin \tilde{\delta} - \sin \hat{\delta} \end{bmatrix} \\ &+ (t - t_{\text{ep}}) \begin{bmatrix} \tilde{\mu}_{\alpha^*} \\ \tilde{\mu}_{\delta} \end{bmatrix} - (t - t_{\text{p}}) \begin{bmatrix} \hat{\mu}_{\alpha^*} \\ \hat{\mu}_{\delta} \end{bmatrix}. \end{aligned} \quad (113)$$

If, on the other hand, the proper motion is not measured relative to the local extragalactic background, but in a global non-rotating frame, then it already includes a contribution from $\boldsymbol{\mu}_0$ and the appropriate observation equations are obtained by deleting the terms depending on \mathbf{a} :

$$\left. \begin{aligned} \tilde{\mathbf{q}}' \boldsymbol{\omega} &= \tilde{\mu}_{\alpha^*} - \hat{\mu}_{\alpha^*} \\ -\tilde{\mathbf{p}}' \boldsymbol{\omega} &= \tilde{\mu}_{\delta} - \hat{\mu}_{\delta} \\ \tilde{\mathbf{q}}' (\boldsymbol{\varepsilon} + (t - t_{\text{fr}}) \boldsymbol{\omega}) &= \Delta\alpha^* \\ -\tilde{\mathbf{p}}' (\boldsymbol{\varepsilon} + (t - t_{\text{fr}}) \boldsymbol{\omega}) &= \Delta\delta \end{aligned} \right\}. \quad (114)$$

6.1.5. The null space vectors

In Sect. 4.4 we introduced the $n \times 6$ matrix \mathbf{V} whose columns span the null space of the normal matrix \mathbf{N} . For completeness we give here the explicit expressions for one such set of null vectors. Any small change in the unknowns \mathbf{x} , by a linear combination of the columns in \mathbf{V} , will leave the calculated residuals unchanged. Applying the frame rotator for arbitrary $\boldsymbol{\varepsilon}$ and $\boldsymbol{\omega}$ obviously leaves the residuals unchanged, and we can therefore compute the columns of \mathbf{V} as the partial derivatives of \mathbf{x} with respect to the six frame rotator parameters. Since we are concerned with infinitesimal changes, the distinction between the AGIS and ICRS frames is no longer necessary. If \mathbf{V} is partitioned similarly to \mathbf{x} and \mathbf{b} in Eq. (31), or $\mathbf{V} = [\mathbf{V}'_s, \mathbf{V}'_a, \mathbf{V}'_c, \mathbf{V}'_g]'$, we find by means of Eq. (114),

$$[\mathbf{V}_s]_i = \begin{bmatrix} q_X & q_Y & q_Z & \tau q_X & \tau q_Y & \tau q_Z \\ -p_X & -p_Y & -p_Z & -\tau p_X & -\tau p_Y & -\tau p_Z \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & q_X & q_Y & q_Z \\ 0 & 0 & 0 & -p_X & -p_Y & -p_Z \end{bmatrix}, \quad (115)$$

where $\tau = t_{\text{ep}} - t_{\text{fr}}$ and we have omitted the source index i on the matrix elements. The order of the astrometric parameters is

$(\alpha^*, \delta, \varpi, \mu_{\alpha^*}, \mu_{\delta})$. From Eqs. (92) and (105) we similarly obtain for the attitude parameters $\mathbf{a}_n = \{a_x, a_y, a_z, a_w\}$,

$$[\mathbf{V}_a]_n = \frac{1}{2} \begin{bmatrix} -a_w & -a_z & a_y & -\tau_n a_w & -\tau_n a_z & \tau_n a_y \\ a_z & -a_w & -a_x & \tau_n a_z & -\tau_n a_w & -\tau_n a_x \\ -a_y & a_x & -a_w & -\tau_n a_y & \tau_n a_x & -\tau_n a_w \\ a_x & a_y & a_z & \tau_n a_x & \tau_n a_y & \tau_n a_z \end{bmatrix}, \quad (116)$$

where $\tau_n = \bar{t}_n - t_{\text{fr}}$. The calibration and global parameters are not affected by the frame rotator, so $\mathbf{V}_c = \mathbf{0}$ and $\mathbf{V}_g = \mathbf{0}$.

Let $\tilde{\mathbf{d}}$ be a vector of small changes to the unknowns \mathbf{x} , as for example the update computed in one of the AGIS iterations. In some situations it is desirable to remove from $\tilde{\mathbf{d}}$ its component in the null space, i.e., to project it on the row space of \mathbf{N} . This will for example ensure that the orientation and spin of the AGIS frame is not, on the average, changed by the update. In principle, we could achieve this by a process analogous to Eq. (87):

$$\mathbf{d} = \tilde{\mathbf{d}} - \mathbf{V}(\mathbf{V}'\mathbf{V})^{-1}\mathbf{V}'\tilde{\mathbf{d}}, \quad (117)$$

where \mathbf{d} is the projection of the update on the row space of \mathbf{N} . This is equivalent to solving the unweighted least-squares problem $\mathbf{V}\mathbf{z} \approx \tilde{\mathbf{d}}$, yielding the orientation and spin components as $\hat{\mathbf{z}} = (\mathbf{V}'\mathbf{V})^{-1}\mathbf{V}'\tilde{\mathbf{d}}$, followed by the subtraction of the null space component $\mathbf{V}\hat{\mathbf{z}}$. In practice, this can equivalently be achieved by means of the frame rotator, without the need to compute \mathbf{V} explicitly.

6.1.6. Role of the frame rotator in AGIS

The frame rotator process consists of the three steps: (i) determine the frame rotator parameters according to Sect. 6.1.4; (ii) correct the astrometric parameters for all sources according to Sect. 6.1.2; (iii) correct the attitude parameters according to Sect. 6.1.3. In principle, this process only needs to be run after convergence of the AGIS iterations; nevertheless, there may be a case for running it after each AGIS iteration, preventing a progressive deviation from the ICRS in the course of the iterations.

In particular, for the simulation experiments described in Sect. 7 it was found necessary to run the frame rotator after each iteration, in order to be able to compare the results of each iteration with the ‘true’ astrometric parameters (used as input to the simulations). Without this correction, the differences between the ‘estimated’ and ‘true’ parameters for the source positions and proper motions would have been grossly distorted by frame errors originating from the starting values of the attitude parameters. In this case all the primary sources were treated as belonging to the subset S_{PM} , for which Eq. (114) is appropriate.

6.2. Selection of primary sources

The astrometric core solution does not use all the sources observed by Gaia, but only a subset of them known as the primary sources. The selection of this subset is made iteratively, based on the results of earlier astrometric solutions and other processes such as double-star and variability analysis (not discussed in this paper). The main criterion for a primary source is that its proper direction, at all times when it is observed by Gaia, is adequately modelled by the standard astrometric model outlined in Sect. 3.2. Examples of sources that should be excluded based on this criterion are solar-system objects, short-period astrometric binaries with a significant size of the photocentre orbit, long-period astrometric binaries with a significant curvature of the photocentre

orbit, certain unresolved binaries where one or both of the components are variable, and active galactic nuclei (AGNs) with significant variation of their photocentre positions. Variable stars, long-period resolved binaries, eclipsing and spectroscopic binaries need not be excluded a priori from the set of primary sources, although many of them are potentially problematic from an astrometric modelling viewpoint (e.g., a variable star might be part of an unresolved double or multiple star, resulting in a variable photocentre position). On the other hand, partially resolved double/multiple stars and other extended sources will be problematic even if their photocentres strictly adhere to the basic astrometric model, Eq. (3), because technically the determination of the photocentre becomes more difficult and less precise.

For the calibrations there are other requirements on the primary sources, in particular that there are enough of them at various magnitudes and colours, while their sky distribution is less important. Securing a sufficient number of primary sources for the calibrations will tend to include many more primary sources in some areas, such as the galactic plane, resulting in a very non-uniform distribution across the celestial sphere.

6.2.1. The number of sources needed for AGIS

The number of sources required for the Astrometric Global Iterative Solution is driven by the calibration needs of having representative numbers of sources of different magnitudes, and the attitude needs of having a sufficient number of sources in every knot interval.

Let us consider first the requirements for the geometric small-scale calibration of the CCDs (Sect. 3.4). The angular extent of a single CCD in the across-scan direction is about 0.1° . It scans the celestial sphere at a rate of 1° min^{-1} and therefore covers about 2000 deg^2 per week, if both fields of view are counted. Since there are 1966 pixel columns across the CCD, we have the convenient rule of thumb that each pixel column covers 1 deg^2 per week. Thus, if the average density of suitable primary sources is $D \text{ deg}^{-2}$ and we require a minimum of N observations of such sources per pixel column for its calibration, then the minimum time needed is N/D weeks. For example, with 10^8 primary sources we have $D = 2400 \text{ deg}^{-2}$, and it is then reasonable that the small-scale calibration can be made, at a resolution of one or a few pixels, in a matter of weeks. However, for the gated observations of bright sources ($G \lesssim 12 \text{ mag}$), the available numbers are much smaller and it will be necessary to sacrifice resolution in time, number of pixels, or both in order to have a sufficient number of observations per calibration cell. For example, the gate used for the brightest sources will perhaps mainly be used for $G \simeq 5.7$ to 8.8 ; the *total* number of such stars is about 130 000 (from the Tycho-2 Catalogue; Høg et al. 2000) or $D \simeq 3 \text{ deg}^{-2}$, of which only a fraction will be suitable as primary sources. Thus, even over the whole five-year mission there will only be a few hundred observations per pixel column. From these considerations it is clear that as many as possible of the bright stars should be selected as primary sources.

Consider next the requirements for the attitude determination. The minimum reasonable knot interval is of the same order as the transit time over a CCD, or about 5 s. Since we require both along-scan and across-scan measurements, and the latter are normally only provided by the SM and some of the AF1 observations, we assume conservatively one observation in each coordinate per field-of-view transit. There are seven CCDs across the width of the field of view; the area scanned is therefore 1000 deg^2 per day in each field of view, or 0.06 deg^2 per 5 s interval. If we require, say, 100 transits per knot interval for a

reliable attitude determination, then the minimum density of primary sources is $D = 1700 \text{ deg}^{-2}$, or some 70 million sources in total. To achieve this density in the galactic pole regions requires that stars as faint as $G \simeq 19$ are included among the primary sources. Thus it is clear that the design aim of $\simeq 10^8$ primary sources is quite reasonable, and that these will have to include both very bright and very faint stars, as well as many quasars down to $G = 20$ for the extragalactic link.

Apart from these minimum requirements, it must be maintained that the quality of the solution will only improve, the more (good) primary sources are included. Stated in the negative sense, the solution quality cannot improve by removing good-quality primary sources. From this viewpoint one should aim to include as many primary sources as possible in the final solution.

On the other hand, one should not forget that it is possible to run AGIS with much fewer primary sources by disabling short-term small-scale calibrations and using longer knot intervals for the attitude. This will increase modelling errors, which however is acceptable for initial runs where the input data have not yet been properly calibrated.

6.2.2. Selection criteria

As outlined above a source has to pass several tests, derived from different processes, in order to qualify as a primary source. The most important test is derived from the AGIS solution itself, and is based on how well the standard astrometric model (Sect. 3.2) fits the data. However, if initially we want to limit the number of primary sources, a somewhat more sophisticated selection procedure is needed to guarantee the minimum requirements.

Each source carries an attribute representing the ‘relegation factor’ U , which is a floating-point number ranging from about 1 to infinity. $U \simeq 1$ implies the source is perfect for use in the astrometric solution, while successively larger values indicate less suitable sources. The relegation factor may incorporate the results of several different tests, and therefore provides a continuous variable for use in the selection process. The name derives from the need to ‘relegate’ a primary source into a secondary one when U exceeds a predefined value, which may happen for example in the course of the AGIS iterations, or from one solution to the next. On the other hand, a secondary source may be promoted to a primary if its U value decreases below the set threshold. This suggests that all potential primary sources should be processed through the source update (Sect. 5.1), after which its status as primary/secondary may be decided.

The excess source noise ϵ_i estimated during the source updating (Sect. 5.1.2) may be a good starting point for calculating the relegation factor, e.g.:

$$U_i = \sqrt{1 + [\epsilon_i/e(G_i)]^2}, \quad (118)$$

where $e(G)$ is a normalization factor depending on the magnitude. The choice of the function $e(G)$ determines the balance between absolute and relative contributions to the modelling error budget. With the choice $e(G) \simeq \sigma_f^{\text{AL}}$ (the formal along-scan observational standard uncertainty for a source of magnitude G), U_i approximates the RMS normalized residual of the source. Selecting sources based on this U_i tends to discriminate against bright stars where modelling errors may dominate over photon-noise errors. On the other hand, choosing $e(G) = \text{constant}$ means that only sources with the smallest ϵ_i are accepted; this may remove too many of the faint sources, where ϵ_i is still small in comparison with σ_f^{AL} . A reasonable compromise between these two

extreme cases should be found. The value from Eq. (118) can later be combined with other factors indicating for example photometric variability, or some other potentially problematic property, so that in general the relegation factor can be determined with some formula using a combination of the source attributes.

Apart from the relegation factor, which indicates whether a source is at all suitable as a primary source, the general principle should be to maximize the total weight of the primary sources. The weight of a source i is defined as

$$W_i = \frac{1}{n_i} \sum_{l \in i} \frac{w_l}{\sigma_l^2 + \epsilon_i^2}, \quad (119)$$

where the average is taken over the $n_i = \sum_{l \in i} 1$ accepted AL observations of the source. Note that we do not use the more obvious definition $W_i = \sum_{l \in i} W_l$ with W_l from Eq. (62), since we do not want to penalize a source by the excess attitude noise in some of its observations, nor favour a source because it is observed many times due to the scanning law.

Having defined the relegation factor and weight per source, the selection of primary sources can be made to maximize the total weight with due regard to sky uniformity (for attitude determination) and magnitude distribution (for instrument calibration). A possible procedure is the following.

We start by specifying the minimum density D_{\min} (deg⁻²) required for the attitude determination, the targeted total number N_{tot} of primary sources, with $N_{\text{tot}} \geq (129600/\pi)D_{\min}$, and the maximum acceptable relegation factor U_{\max} . For the geometric calibration of gated observations we may also specify minimum numbers $\{N_g\}$ for several intervals in G . Then:

1. Using a coarse-grained tessellation of the celestial sphere (see below), select in each pixel the N_p sources with the largest W_i that satisfy $U_i \leq U_{\max}$, where N_p is the minimum number per pixel that will ensure D_{\min} .
2. For each magnitude bin with a required minimum number N_g , count the actual number of primary sources already selected; if it is less than N_g add sources with $U_i \leq U_{\max}$ based on W_i .
3. If the total number after Step 1 and 2 is less than N_{tot} , add sources with $U_i \leq U_{\max}$ based on W_i . If the total number exceeds N_{tot} , reduce U_{\max} and repeat the process.

If the required number cannot be reached in a particular pixel or magnitude bin, then it is necessary to increase U_{\max} locally for that pixel or bin.

For the tessellation in Step 1 in principle any reasonable way to divide up the sphere into cells of approximately the same area could be used, but for statistical operations the HEALPix scheme (Górski et al. 2005) has some advantages (O’Mullane et al. 2001). The cell size is determined by the choice of HEALPix parameter $NSIDE$ and is important as it predicts the level of homogeneity over the sphere. A cell area of about 1/3 of the field of view might be close to optimal, and is achieved with $NSIDE = 128$ yielding 196 608 cells.

For the first run with real data some selection must be made using the initial star catalogue. In this case the relegation factor may be set to 1 for all sources and the selection based entirely on their spatial distribution and magnitudes. This will reduce the input to the first run of the astrometric solution, after which the relegation factor will be updated as described above. The secondary-source update step runs on all sources not included in AGIS; hence this will set a relegation factor for all sources observed by Gaia.

6.3. Computation of standard uncertainties and correlations

It is mandatory that the catalogue of astrometric parameters resulting from the astrometric core solution includes complete and reliable information about the expected error statistics. The most important quantity is the estimated standard uncertainty of each astrometric parameter. However, the statistical correlation between the different astrometric parameters – both between the different parameters of the same source and between the parameters of different sources – is also important and should be quantified. Such correlations are produced both by attitude modelling errors (Sect. 5.2.5) and the statistical uncertainty due to the finite astrometric weight of the sources contributing to the attitude determination. More generally, we need a method to estimate the 5×5 covariance matrix $\text{Cov}(s_i, s_j)$ of the astrometric parameters of any two sources i, j (including the case $i = j$). In principle, these are sub-matrices of the upper-left $n_s \times n_s$ part of N^\dagger , the pseudo-inverse of the complete normal equations matrix in Eq. (30). (The pseudo-inverse should be used since the matrix is singular.) Although there are methods to compute selected elements in N^\dagger that may be feasible even for a system as large as this, it is utterly impossible to produce any significant fraction of the covariances by a direct computation. Instead, it will be necessary to rely on approximations and statistical estimates. A first approximation is obtained by ignoring the statistical uncertainty contributed by the errors of the attitude, calibration and global parameters; in this case we can ignore all parts of N in Eq. (31) except N_{ss} and find

$$\text{Cov}(s_i, s_j) \equiv [N^\dagger]_{ss} \approx \begin{cases} [N_{ss}]_{ii}^{-1} & \text{if } i = j, \\ \mathbf{0} & \text{if } i \neq j, \end{cases} \quad (120)$$

where the inverse of $[N_{ss}]_{ii} = A_i^\dagger W_i A_i$ is regularly computed as part of the source updating using Eq. (57). However, in this approximation we clearly cannot estimate the covariance between sources ($i \neq j$), which is unacceptable; moreover, we underestimate the within-source covariance ($i = j$) because of the neglected attitude and calibration errors. Refining this estimate is a very important problem which however is addressed elsewhere (Holl et al. 2010, 2012).

A somewhat related problem is the need to be able to transform the astrometric results, without loss of information, to an arbitrary epoch different from the t_{ep} used in the astrometric solution. The standard model in Eq. (4) allows the astrometric parameters to be transformed in a completely reversible manner, based on the assumption of uniform space motion relative to the solar-system barycentre. In this process it is necessary to include the sixth astrometric parameter μ_r , even if it is (partially) derived from a spectroscopic radial velocity. Similarly, the transformation of the covariance matrix must consider all six parameters. The relevant formulae are given in Vol. 1, Sect. 1.5.5 of The Hipparcos and Tycho Catalogues (ESA 1997) and are not repeated here.¹⁷ Indeed, the normal equations for the six parameters contain the full information of the Gaia observations of a particular source with respect to the standard astrometric model, and for this reason it is desirable to compute and store these normals even if only a subset of them is used in the actual source update (cf. Sect. 5.1.3).

¹⁷ The sixth parameter μ_r is denoted ζ in ESA (1997).




Fig. 8. Simplified architectural design diagram of AGIS. The rounded boxes are independent Java processes running in parallel on different nodes of a multi-CPU processing cluster. The arrows indicate main data exchanges between the various processes. Input/output-related data flow from/to the storage system is not shown, and likewise some important but conceptually irrelevant interactions between some elements (e.g., the *Servers* and the *RunManager*).

7. Software implementation and demonstration solutions

The practical realisation of the AGIS scheme as outlined in the preceding sections is contained in software that is being developed, since early 2006, jointly by teams at the European Space Astronomy Centre (ESAC) and Lund Observatory within the Gaia Data Processing and Analysis Consortium (DPAC). It is a central software module embedded into a complex, overall data processing system (O’Mullane et al. 2007), whose ultimate goal is the creation of the final Gaia catalogue with a targeted release date around 2021.

This section gives a concise overview of the architectural design of the AGIS software (Sect. 7.1) on the one hand, and on the other presents some selected results from a recently (June–November 2011) conducted large-scale astrometric solution (Sect. 7.2) using as input simulated data for more than 2 million sources. While this number is still a factor 50 smaller than the number of primary sources foreseen in the final AGIS runs (around 2018–2020), and the present simulations are much simplified especially with respect to the attitude modelling, we believe that this proof-of-concept run demonstrates the practical validity and correctness of the key theoretical concepts described in this paper. Future tests will involve simulated data sets that are both larger and more realistic, with a parallel further development of the algorithms and software in view of the added new complexities.

7.1. AGIS software overview

The term AGIS is subsequently often used to refer to the actual software implementation of the scheme. Like virtually all Gaia data processing software, AGIS is entirely written in the object-oriented Java programming language (O’Mullane et al. 2010).

The implementation has been briefly outlined in Lammers et al. (2009) and more comprehensively in O’Mullane et al. (2011), to which the interested reader is referred for more details. In this section, some key classes are briefly described; following Java naming conventions their names (given in *italics*) are concatenated capitalized nouns, as in *RunManager*.

Owing to the number of sources and the associated large data volumes that have to be handled (see Sect. 1) it is clear that a well-performing system must be distributable on modern, multi-node, multi-core processing hardware environments and make optimal use of parallelism as far as permitted by the AGIS scheme. Another elementary consideration is that inherently slow disk input-output operations should be minimized and never allowed to be a bottleneck for any of the computing processes.

These basic requirements have led to the system schematically depicted in Fig. 8 with its main components and data flow. The central elements are *DataTrains*, *Servers*, a *RunManager*, and a *ConvergenceMonitor*. When an AGIS run starts the *RunManager* splits the entire processing task of the first iteration into separate and independent jobs which are then taken and executed in parallel by *DataTrains* that have been started on the different CPUs of the processing system. Each such job involves the processing of all observations for a group of sources (with typically 100–1000 sources per group), which is done by looping over the sources, one at a time. Before the loop is entered all the data that are needed for the processing (observations, as well as all source, attitude, and calibration data) have been loaded into memory and are passed on to the core algorithms. This is a key design aspect which, together with a suitable grouping of the data on the storage system, ensures that the system is never input/output limited and that it has a constantly high utilization of the CPUs (typically at the 90% level).

Each AGIS iteration starts with the updating of the respective sources (see Sect. 5.1). Computed provisional updates are then written to the storage system and, finally, the observations together with the updated source data are sent to attitude, calibration, and global servers via the CPU-interconnecting network. The servers themselves are distributed in different ways (see, e.g., Sect. 5.2.2 for attitude), but all are similar in that they accumulate normal equations by adding observation equations as outlined in Sects. 5.2.1, 5.3, and 5.4, respectively. When all *DataTrain* jobs are finished, the *RunManager* signals to the servers that all observations have been sent. This triggers that a Cholesky decomposition (Appendix C) is made of the accumulated normals matrices in every server, that the partial results on different servers are combined where necessary (e.g., the attitude segments are joined), and finally that the results are persisted in the storage system. That marks the end of an iteration. Subsequent iterations are then started in the same manner as the first one, viz., through the creation of a set of processing jobs.

Iteration k uses the computed updates from iteration $k-1$ and generates updated parameters for use in the following iteration $k+1$. This progress is monitored by the *ConvergenceMonitor* through the accumulation of a selected list of statistical quantities in the form of graphical plots (e.g., histograms of the updates of all the astrometric parameters) accessible in real time through a web interface. Naively, one may expect that the system can be considered converged if the updates become smaller than a pre-defined limit; however, finding an unambiguous and automatically verifiable convergence criterion has proven to be a surprisingly complex problem (see Sect. 4.4 in Bombrun et al. 2011). We believe now that human inspection is indispensable

Table 1. Characteristics of the simulated input data and demonstration solution.

Quantity	Value
Duration of science mission	5.0 yr
Number of sources	2 256 222
Number of along-scan (AL) observations	1.625×10^9
Number of across-scan (AC) observations	1.805×10^8
Standard uncertainty per AL / AC observation (representative):	
$G \leq 13$	$92 \mu\text{as} / 520 \mu\text{as}$
$G = 15$	$230 \mu\text{as} / 1350 \mu\text{as}$
$G = 17$	$590 \mu\text{as} / 4000 \mu\text{as}$
$G = 18$	$960 \mu\text{as} / 7600 \mu\text{as}$
$G = 19$	$1600 \mu\text{as} / 16000 \mu\text{as}$
$G = 20$	$2900 \mu\text{as} / 38000 \mu\text{as}$
Number of astrometric parameters	1.128×10^7
Number of attitude spline knots	6.575×10^5
Number of attitude parameters	2.630×10^6
Number of calibration parameters	7 812
Number of global parameters	1

to assess the convergence status of the system reliably and, ultimately, decide on the termination of the iterative loop.

An important feature of the *RunManager* is the ability to use different algebraic solution methods by selecting among different available iteration schemes. The description above explains the ‘kernel’ computation of provisional updates to the unknowns employing the Gauss-Seidel-type preconditioner approximation (Eq. 53 in Sect. 4.6) to the full normal matrix of the system (Sect. 4.5). How these provisional updates are actually combined at the end of an iteration to form the final updates to the unknowns depends on the chosen iteration scheme. AGIS can use all four schemes outlined in Sect. 4.7, viz., the simple iteration (SI), accelerated simple iteration (ASI), conjugate gradients (CG), and hybrid scheme (A/SI-CG). The previous description essentially refers to SI; in the other schemes there are a few extra steps which however are immaterial for understanding the software system.

AGIS is controlled through a list of a few hundred key-value parameters (‘properties’) which are configured before a run starts. Examples are: the numbers of servers and threads, the size of *DataTrain* jobs, the starting values for the unknowns, and the employed solution method. Also, which update blocks are active during a run is controlled via properties. Any combination that involves at least a source update is possible, e.g., SACG, SA, SCG.

The optimum number of data trains in a run is a complex trade-off between the available number of CPUs and memory, usable network bandwidth (more trains create more inter-CPU traffic), and the given maximum storage system throughput. By design of the system, the run time should scale inversely with the number of data trains (assuming there are enough CPUs), i.e., doubling the number of *DataTrains* should halve the run time. In the tests done until now, the run times are very satisfactory; however, more work remains for achieving the desired optimal scaling behaviour. A first AGIS run using simulated data for a 5 yr mission with 50 million primary sources was successfully completed in June 2011. This being only a factor 2 less than the baseline 100 million sources envisaged in the final AGIS run towards the end of the mission, it marks an important milestone in the development of the operational system.

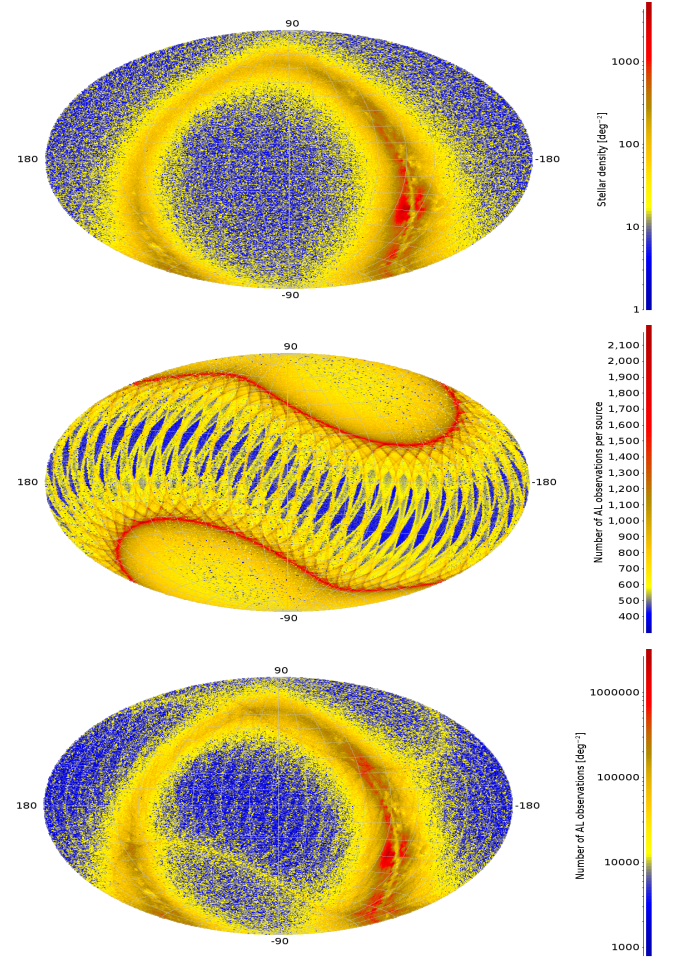


Fig. 9. All-sky projections of (from top to bottom) the total stellar density in the input data to the demonstration solution, the number of AL observations per source, and the resulting spatial density of AL observations. These and all subsequent sky maps use the Hammer–Aitoff equal-area projection in equatorial coordinates, with α running from -180° to $+180^\circ$ right to left. **Top:** The simulated sky contains some 2 million single stars covering the Gaia magnitude range $6 \leq G \leq 20$. The density ranges from less than 1 deg^{-2} around the galactic poles to a maximum of about 4800 deg^{-2} near the galactic centre in the bottom-right quadrant of the map. **Middle:** The number of along-scan observations per source reflects the scanning law of Gaia, which is roughly symmetric around the ecliptic plane and gives an overabundance of observations at ecliptic latitudes $\pm 45^\circ$. **Bottom:** The combination of the source density and the scanning law gives the displayed density of along-scan observations.

7.2. Demonstration solution

7.2.1. Data simulation and model assumptions

Since the start of the development in early 2006, AGIS has been tested continuously using simulated datasets of varying complexity and size generated by the Gaia System Simulator (Masana et al. 2005) created by DPAC’s dedicated coordination unit for Data Simulations (CU2; Mignard et al. 2008; Luri & Babusiaux 2011). In the following we present the results of a test solution using 5 years of simulated astrometric observations for about 2 million well-behaved (single) stars with a realistic distribution both in magnitude and coordinates, based on




Fig. 10. Map of the parallax errors in iteration 1. The iterative astrometric solution starts off with spatially correlated errors in the astrometric parameters, generated as described in the text. These (initial) parallax errors have amplitudes of a few tens of mas. The number at the top-right corner of this (and subsequent) maps is the maximum value of the displayed range in μas .

the Besançon galaxy model (Robin et al. 2003). Figure 9 (top) shows the spatial source density of the data set in equatorial coordinates. Of particular interest for the AGIS run is the stark density contrast between ~ 1 and 5000 deg^{-2} mainly depending on galactic latitude, resulting in similarly high ratios in the total astrometric weight of the sources in Gaia’s two fields of view. In Bombrun et al. (2011) it was shown (using simulations on a smaller scale) that a high weight contrast tends to reduce the convergence rate of the astrometric solution compared to a situation where the weights are more balanced; however, the solution always converges to the correct solution provided that enough CG iterations are carried out. We will show that this key result is confirmed in the demonstration solution.

The input data were generated using a fully-relativistic model of the observed (proper) directions $\mathbf{u}_i(t)$ in Eq. (7), including gravitational light deflection for PPN parameter $\gamma = 1$, assuming the Nominal Scanning Law (Sect. 3.3), the nominal geometrical instrument model (Sect. 3.4) and nominal performance of the instrument (in particular the centroiding accuracies σ_l^{AL} , σ_l^{AC} as functions of G ; see Table 1). However, in order to test the capability to recover a varying basic angle, a step-wise perturbation was introduced corresponding to the sinusoidal variation of $\Delta\Gamma_j$ in Eq. (19) with a period of 2.5 yr and an (unrealistically large) amplitude of $500 \mu\text{as}$. In the astrometric solution, the large-scale AL calibration interval was set to 30 days, matching the step width of the perturbation signal. The solution used only one global parameter, viz., $g_0 = \gamma - 1$ as described in Sect. 5.4. Table 1 lists some statistics of the data and solution, while the middle and bottom maps in Fig. 9 show how the number of observations varies across the sky.

Not all elements of the numerical algorithms described in this paper have as yet been integrated into the running software system which otherwise implements the basic model described in Sect. 3. In particular, the estimation of source and attitude excess noise (Sects. 5.1.2 and 5.2.5) were not activated in the present solution. This was not a problem for the demonstration run, as the applied observation noise is well-behaved, purely Gaussian. Further development of the AGIS software will to a large extent focus on making the solution robust against all kinds of unexpected input data.

The demonstration solution was run with all four update blocks (S, A, C, G) enabled, using starting values for the attitude, calibration and global parameters that were erroneous on the mas level. (As explained in Sect. 4.5 and footnote 9, the re-

sults of the subsequent iterations are independent of the initial values for the source parameters, because the updating always starts with the source block.) These initial values were created by adding Gaussian, uncorrelated errors to the true attitude parameters, with a standard deviation of 50 mas , using the nominal calibration parameters (i.e., excluding the sinusoidal modulation of the basic angle), and a value of $g_0 \equiv \gamma - 1 = 0.1$ (Sect. 5.4). An attitude knot interval of 240 s was used in order to have a sufficient number of observations per degree of freedom even at the galactic poles. This interval is short enough that the attitude splines are able to represent the true attitude (i.e., the analytical nominal scanning law) with an RMS error of less than $9 \mu\text{as}$. Although this is larger than the modelling errors aimed at in the real data analysis, it is sufficiently small in comparison with the typical attitude estimation errors ($\geq 20 \mu\text{as}$; see Fig. 11) to have a negligible impact on the overall astrometric accuracy of the present solution.

During the source update in the very first iteration, the initial errors in the attitude, calibration parameters and γ propagate to the sources, creating astrometric errors of a few tens of mas (Fig. 10) that are spatially correlated on a scale comparable to the attitude knot interval ($\sim 4^\circ$). These errors are quite hard to remove in subsequent iterations, but may be representative of the situation encountered by AGIS when processing the real mission data based on a fairly uncertain initial attitude and instrument calibration.

With these starting conditions, 135 iterations were carried out using the conjugate gradients (CG) scheme. A re-initialisation of the CG scheme was made after the first 40 iterations to avoid the development of a much slower convergence phase observed in some previous runs; after this, no further re-initialisation was made. At iteration 135 the typical updates of, for example, the parallaxes and the along-scan attitude were at or below a level of $5 \times 10^{-4} \mu\text{as}$ (Fig. 11). This is still slightly above the numerical noise floor set by the double-precision arithmetic ($\sim 10^{-16} \text{ rad}$ or $\sim 10^{-5} \mu\text{as}$). As discussed by Bombrun et al. (2011), truncating the iterations before the numerical noise floor has been reached implies the presence of spatially correlated ‘truncation errors’ having an amplitude of a few times the typical updates, or $\sim 10^{-3} \mu\text{as}$ in the present case. We now proceed with a more detailed analysis of the results.

7.2.2. Source results

In this section we focus on the results for the parallax (ϖ), which is arguably the most interesting parameter from an astrophysical viewpoint; moreover, as a scalar quantity independent of epoch and reference frame, its statistics can be summarized compactly and without ambiguity. However, the behaviour of the position and proper motion parameters is qualitatively similar, and some results are given in Table 2 and Fig. 14.

The top diagram in Fig. 11 shows the typical sizes of the errors and updates in parallax versus iteration number, as measured by the Robust Scatter Estimate¹⁸ (RSE). The parallax errors settle relatively quickly (around iteration 25) at a level of $146 \mu\text{as}$ and remain stable till the end of the run with updates becoming successively smaller, reaching the level of $10^{-3} \mu\text{as}$ around iteration 120. The actual error distribution is symmetric but strongly non-Gaussian (in fact more like a Laplace distribution) due to

¹⁸ The RSE is defined as 0.390152 times the difference between the 90th and 10th percentiles of the distribution of the variable. For a Gaussian distribution it equals the standard deviation. The RSE is used as a standardized, robust measure of dispersion in CU3.




Fig. 11. Top: Evolution of the typical parallax error (crosses) and parallax update (circles) as functions of the iteration number. The typical error settles at around $146 \mu\text{as}$. **Bottom:** Evolution of the typical attitude error (crosses) and update (circles) as functions of the iteration number, for the three principal SRS axes x (red), y (blue), and z (black). The errors settle at around 167, 224, and $20 \mu\text{as}$, respectively. In both these plots the typical errors and updates are given by the Robust Scatter Estimate (RSE), similar to an RMS value (see footnote 18).

the variation of star numbers and observational standard uncertainties with magnitude (cf. Tables 1 and 2). The overall parallax error RSE of $\sim 150 \mu\text{as}$ is therefore representative for the median magnitude, $G \simeq 19$, in good agreement with current accuracy predictions (Lindegren 2010). The overall sizes of the errors and updates shown in Fig. 11 should however be complemented by more detailed statistics as functions of coordinate and magnitude.

Figure 12 shows the spatial distribution of the parallax errors at a few selected iterations. The left column shows the median error in each cell, while the right column shows the median absolute value of the error. These quantities serve as robust proxies for the mean and RMS values (the RSE is not used for the latter as many cells have too few sources for this measure), and therefore may suggest the levels of systematic and random errors as function of position. Figure 13 shows the corresponding maps for the parallax updates. After a few iterations, when the overall parallax errors are already below the 1 mas level, very significant systematic (i.e., spatially correlated) errors of a few mas remain, especially in the high-density areas of the galactic plane. These are damped in the subsequent iterations, but still remain at a level of several hundred μas in the galactic centre region around iteration 20, when the overall parallax errors start

to settle at their final value according to Fig. 11. At iteration 135 the regional errors have virtually disappeared, and the error map shows a characteristic pattern with the solution being seemingly better around the galactic equator than in the polar regions. This is purely an effect of number statistics: in the galactic pole areas the cells contain rather few sources (often just a single source, in which case the displayed values are simply the individual parallax errors), while closer to the galactic plane the scatter from one cell to the next is reduced by the median-averaging over many sources.

The sequence of error maps for iterations 5, 20, and 135 corroborates a key result of Bombrun et al. (2011), viz., that by iterating long enough the system can cope with large spatial imbalances in the astrometric weights, provided that the minimum source density allows a good attitude determination at all points.

The median absolute values of the parallax errors shown in the right column of Fig. 12 quickly settle in a large-scale pattern that mainly reflects the expected variation of parallax accuracy with ecliptic latitude (see, e.g., Table 3 in Lindegren 2010). This, in turn, depends on the scanning law, i.e., on a combination of the number of observations per source (see Fig. 9, middle) and the geometric configuration of the scans – for example, the over-density of observations at $\pm 45^\circ$ ecliptic latitude does little to improve the parallaxes, which are then mainly producing shifts in the AC direction, while Gaia is primarily sensitive to the AL displacement. Number statistics reduce the between-cell scatter of the median absolute values as well, which accounts for the smoother appearance along the galactic equator.

The update maps in Fig. 13 conform with expectations and the preceding discussion. Of particular interest from a diagnostic viewpoint is the observation that the amplitude and spatial distribution of the median updates in the non-converged solution give a fair indication of the (systematic) truncation errors. This is obviously useful for assessing the state of convergence, as the update maps can be constructed for the real mission data as well (whereas the error maps are of course unknown). For example, based on the median updates in iteration 20 (middle left map of Fig. 13) one might correctly conclude that truncation errors of a few hundred μas remain, especially in the galactic centre region, as shown in the middle left map of Fig. 12. By the same reasoning it appears that truncation errors at iteration 135 should be well below $1 \mu\text{as}$.

In tests using other datasets with lower density contrasts we have seen a more rapid convergence. A similar slowdown in convergence for a non-uniform weight distribution was observed and discussed in Bombrun et al. (2011) based on small-scale simulations. A likely mechanism for this slowdown is related to the weight contrast problem discussed by van Leeuwen (2005) in connection with the new reduction of the Hipparcos data. Any of the fields of view scanning through a high-density area (or an area with many bright stars) creates a strong astrometric weight imbalance between the two viewing directions, as the other field usually points to an area of the sky with much lower source density (or fainter stars). Errors in the along-scan attitude create correlated errors in the parameters of all sources observed at that time, whether they are in the preceding or following field of view. If these source parameters are then used to correct the attitude, with little counterbalancing effect from the (relatively few) sources in the other field, the attitude may get only marginally improved, with the net effect of slowing down the damping and de-correlation of the errors in the high-density areas. Thus, more




Fig. 12. Maps (in equatorial coordinates) of the parallax errors in the three selected iterations $k = 5, 20, 135$ (top to bottom). The left column shows the median of the parallax error $\varpi^{(k)} - \varpi^{\text{true}}$, while the right column shows the median of the absolute error $|\varpi^{(k)} - \varpi^{\text{true}}|$; each map cell (of about 0.84 deg^2) shows the colour-coded value computed for the sources located in that cell. Empty cells are shown in black. On every map plot the top left label indicates the iteration number and the top right label is the maximum value of the displayed range in μas . See text for further details.

iterations are needed compared to a more weight-balanced situation.¹⁹

Figure 14 shows the error and update maps of the proper motions at iteration 135. Since proper motions are vectors, the displayed quantities are the median lengths of the vectorial differences, which are non-negative by definition. The maps are in expected agreement with the corresponding parallax ones concerning visible patterns and structures. A prominent feature absent in the case of the parallaxes is the lighter bands of relatively smaller proper motion errors around ecliptic latitudes $\pm 45^\circ$ caused by the oversampling of these parallels by the scanning law (cf. Fig. 9). In contrast to the parallax case, these observations do contribute to the determination of the proper motion, especially for the component in ecliptic longitude.

All spatial maps in Figs. 12–14 show median values computed from distributions with stars of all magnitudes. A lot more information is contained in the magnitude-resolved versions of

the maps, which are not presented here for brevity. Instead, Table 2 shows the RSEs (see footnote 18) of the errors and normalized errors (see below) of all the astrometric parameters in iteration 135, subdivided according to magnitude. The normalized error is defined as the error (in the case of right ascension, $\Delta\alpha^* \equiv \Delta\alpha \cos \delta$) divided by the corresponding formal standard uncertainty obtained from the inverse of the source normal matrix, i.e., by using the approximation in Eq. (120). As discussed in Sect. 6.3 this approximation underestimates the true standard uncertainties of the astrometric parameters by neglecting the contribution from the attitude uncertainty, which may be particularly important for the bright stars where the photon noise is relatively less important.

The RSE of the errors in Table 2 are in reasonable agreement with recent mission accuracy assessments. For example, compared with Eqs. (5.2)–(5.5) in Lindegren (2010) the present values are a few per cent larger for $G < 15$, and up to some 20% smaller for the fainter stars. This suggests a more conservative photon-statistical error budget in Lindegren (2010), combined with a larger-than-nominal contribution from the attitude uncertainty in the present solution. The latter effect, further discussed in Sect. 7.2.3, is indeed to be expected in the present demonstra-

¹⁹ In the new reduction of the Hipparcos data by van Leeuwen (2007) the weight ratio was artificially damped in order to improve the connectivity between the two fields of view in high-contrast situations. We have found that this is not required for Gaia provided that the solution is iterated to convergence; see Bombrun et al. (2011).




Fig. 13. Same as Fig. 12 but showing the updates in parallax, i.e., the median values of $\varpi^{(k)} - \varpi^{(k-1)}$ (left column) and $|\varpi^{(k)} - \varpi^{(k-1)}|$ (right column).

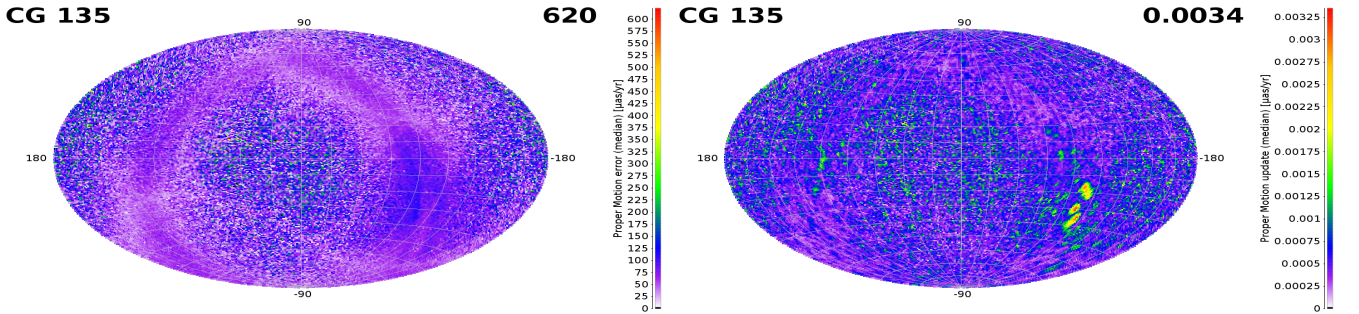


Fig. 14. Maps of the error (left) and update (right) in proper motion for iteration 135. Each cell shows the colour-coded median error/update in units of $\mu\text{as yr}^{-1}$, where the individual error/update is computed in terms of the equatorial components as $(\Delta\mu_{\alpha*}^2 + \Delta\mu_{\delta}^2)^{1/2}$.

tion solution using far fewer primary stars than planned for the real mission.

A more stringent test of the quality of the solution is obtained by considering the RSE of the normalized errors (i.e., after division by the formal standard uncertainties as described above), which is here denoted ϱ . Ideally we should have $\varrho \simeq 1$ for any parameter and any magnitude. As shown in Table 2, this is very nearly the case in all parameters for $G > 15$, meaning that the actual errors are roughly consistent with the standard uncertain-

ties computed from Eq. (120). For the brighter stars ϱ becomes progressively larger, supporting the interpretation that the attitude uncertainty has a significant impact on the accuracy of the bright stars in this solution. For example, the value $\varrho = 1.369$ obtained for the parallaxes of stars brighter than $G = 13$ suggests a quadratic attitude contribution to the parallax errors for these stars of $[7.5^2 - (7.5/1.369)^2]^{1/2} = 5.1 \mu\text{as}$, while a similar computation for the next three magnitude bins gives 5.7, 5.1 and 5.4 μas (with a rapidly increasing uncertainty). Thus it appears that the

Table 2. RSE of the errors in the astrometric parameters and of the normalized errors (i.e., after division by the formal standard deviations) for different magnitude ranges. The asterisk in α^* indicates that the errors are true arcs on the sky; cf. footnote 11. See text for further explanations.

Magnitude range	No. stars	RSE of error [μas and $\mu\text{as yr}^{-1}$]					ϱ = RSE of normalized error [unitless]				
		α^*	δ	ϖ	μ_{α^*}	μ_{δ}	α^*	δ	ϖ	μ_{α^*}	μ_{δ}
$G < 13$	18 253	6.6	5.7	7.5	4.5	4.0	1.507	1.470	1.369	1.443	1.425
$13 \leq G < 15$	70 355	12.4	10.6	14.9	8.7	7.5	1.112	1.100	1.082	1.098	1.100
$15 \leq G < 16$	88 116	20.2	17.3	24.9	14.3	12.3	1.024	1.024	1.022	1.032	1.024
$16 \leq G < 17$	151 639	30.8	26.7	38.4	21.8	19.0	1.010	1.010	1.010	1.014	1.008
$17 \leq G < 18$	272 424	49.4	42.8	61.8	34.8	30.4	1.004	1.006	1.003	1.002	1.003
$18 \leq G < 19$	489 253	83.3	70.7	104.1	58.9	50.8	1.001	1.001	1.002	1.003	1.002
$19 \leq G$	1 166 182	167.9	140.0	207.6	118.5	100.2	1.001	1.000	1.000	1.001	1.000
all G	2 256 222	116.8	98.7	145.6	82.4	70.6	1.009	1.009	1.007	1.009	1.008

values $\varrho > 1$ found for the parallaxes of the brighter stars can be accounted for by assuming a constant contribution, by about 5–6 μas RMS, to the parallax errors from attitude and/or calibration errors. For the other astrometric parameters we similarly find a constant RMS contribution to the positional errors of about 4–5 μas , and to the proper motion errors of about 3–4 $\mu\text{as yr}^{-1}$. It will be shown below that these numbers are consistent with the actual attitude errors found in the solution.

It should be noted that the reference epoch for the astrometric parameters was set to exactly half-way into the mission, i.e., $t_e = 2014.5$ for the simulated mission interval 2012.0–2017.0. This is optimal in the sense that the positional uncertainties are minimized for approximately this epoch, and that the errors in position and proper motion are nearly uncorrelated. A reference epoch half-way through the mission was also assumed for the accuracy assessment in Lindegren (2010), with which the present results have been compared.

In this solution, the median value of the parallax errors of all the 2.2 million sources was not significantly different from zero (the actual value was +0.004 μas). This shows that, in the absence of systematic observational errors, the astrometric solution is able to determine the *absolute* parallaxes of the sources, as could be expected from the basic principles of the mission. It is especially worth noting that this was achieved while simultaneously determining the PPN γ parameter, known to be strongly correlated with the parallaxes (cf. Sect. 5.4).

7.2.3. Attitude results

The bottom diagram in Fig. 11 shows the RSE attitude errors and updates as a functions of the iteration number, where the small differences of the attitude quaternions have been transformed into small rotations along the SRS axes (x , y , z) according to Sect. A.6 and expressed in μas . The error component around the z axis corresponds to the AL attitude error, while the x and y components are linear combinations of the AC attitude errors in the PFoV and FFoV.²⁰ The z (AL) errors settle at an overall level of 20 μas around iteration 60, while the updates continue to decrease in a similar manner as for the parallaxes (Fig. 11, top). The RSE values of the x and y errors converge to 167 μas and 224 μas , i.e., an order of magnitude larger than in z , reflecting the larger observational errors in the AC direction (Table 1) and

the smaller number of AC observations. The ratio of the errors about y and x , $224/167 \approx 1.34$ is in perfect agreement with the value expected from the geometry of the observations (Fig. 2), viz., $\tan(\Gamma_c/2) \approx 1.34$ for a basic angle of $\Gamma_c = 106.5^\circ$.

The converged AL attitude error of 20 μas is completely consistent with the previously inferred constant RMS contribution, by 5–6 μas , to the parallax errors (see Sect. 7.2.2), as can be seen from the following considerations. For most stars, the propagation of random observational errors from individual AL observations to the parallaxes (say) is largely governed by geometrical factors and the total number of observations per star, and can be statistically described by a ‘coefficient of improvement’ which can be estimated to $207.6/2300 \approx 0.09$ using the RSE of the parallax errors for the faintest bin in Table 2 combined with the typical AL observational error at $G \approx 19.6$ (cf. Table 1). This factor assumes that the individual observational errors (at each CCD) are uncorrelated, which is a very good approximation for the photon-statistical centroiding errors, but not for the attitude errors, which have a correlation length determined by the knot interval of the attitude spline. In the demonstration solution, the knot interval was 240 s, which is much longer than the time it takes an image to cross the nine CCDs in the astrometric field (≈ 40 s). Therefore it is a much better approximation to assume a *constant* attitude error for the whole field crossing, corresponding to nine AL observations. As a result, the coefficient of improvement relevant for the attitude error should be a factor three larger, or ≈ 0.27 . The AL attitude uncertainty of 20 μas therefore corresponds to $20 \times 0.27 \approx 5.4$ μas in the parallax, in very good agreement with the empirical result of 5–6 μas . For the other astrometric parameters a corresponding calculation yields a contribution of about 4 μas in position and 3 $\mu\text{as yr}^{-1}$ in proper motion. Although these numbers are somewhat smaller than the empirical estimates in Sect. 7.2.2 (possibly indicating an additional contribution from the calibration errors), the overall agreement is striking.

The attitude errors obtained in the solution ultimately come from the observational errors of the individual observations, through the process of fitting the attitude spline functions to these observations. If more observations (of the same quality) are added, the attitude errors are expected to diminish inversely with the square root of the number of observations (as long as the modelling errors are not a limiting factor). The present AL attitude error of 20 μas is roughly what can be expected from the density and magnitudes of primary sources in the demonstration solution, as can be seen from the following simple calculation. The AL attitude has essentially one degree of freedom per knot

²⁰ The attitude errors discussed here must not be confused with the excess attitude noise ϵ_a introduced in Sect. 3.6. The latter represents modelling errors, which are practically absent in the demonstration solution.

interval (240 s). The average number of AL observations per degree of freedom is therefore about 2500 (Table 1). From the magnitude distribution in Table 2 and the AL observational uncertainties in Table 1 one can estimate that the average AL observation carries a statistical weight (σ_{AL}^{-2}) corresponding to an AL standard uncertainty of $\sigma_{\text{AL}} \approx 650 \mu\text{as}$ (taking into account the weight reduction by ϱ^{-2} for the bright stars). The mean flow of observations therefore should allow the AL attitude to be pinned down with an uncertainty of about $650 \times 2500^{-1/2} \approx 13 \mu\text{as}$. However, this rough calculation assumes uniform distribution of the stars across the sky, with the same mean density (55 deg^{-2}) as in the demonstration run. Considering that large parts of the sky have a much lower density (typically $5\text{--}10 \text{ deg}^{-2}$; see Fig. 9, top), which implies a less precise attitude at the corresponding times, and that we have also neglected the attitude modelling errors, which here amount to at most $9 \mu\text{as}$ RMS (Sect. 7.2.1), it is not unreasonable that the overall AL attitude uncertainty is about 50% larger than according to the above calculation.

The demonstration run uses only 2% of the stars envisaged for the final AGIS solution, and the majority of them are faint, whereas the real primary stars are preferentially selected among the brighter stars when possible (cf. the discussion in Sect. 6.2.2). For a model distribution of 10^8 primary sources similar to the one described by Hobbs et al. (2010), the AL observational uncertainty corresponding to the average statistical weight is more like $200 \mu\text{as}$, rather than the $650 \mu\text{as}$ in the present data. On the other hand, the attitude knot interval will also be much shorter than the 240 s used in the present run, perhaps even as short as 5 s, which is about the shortest knot interval that can reasonably be used in view of the normal CCD integration time of 4.42 s. Combining these numbers we estimate that the final AGIS run on the real Gaia data might obtain an AL attitude uncertainty, due to the photon noise, of about $(20 \mu\text{as}) \times (200/650) \times [0.02 \times (240/15)]^{1/2} \approx 6 \mu\text{as}$. To this should be added the attitude modelling error (i.e., how well the spline can represent the effective attitude), which is difficult to estimate without more reliable information about attitude irregularities (Appendix D.4). Generally speaking, the optimum knot interval will roughly balance the estimation and modelling uncertainties, so that the total uncertainty is less than twice the estimation uncertainty. Assuming a total AL attitude error of $12 \mu\text{as}$ RMS, this would give less than $4 \mu\text{as}$ RMS to be added quadratically to the parallax uncertainties. Thus, the attitude contribution to the final astrometric parameters appears to be relatively small even for the bright stars. However, this does not take into account the additional complications caused by the gated observations (Appendix D.3) and residual calibration errors due to, for example, CTI effects (Appendix D.2).

7.2.4. Calibration results

The calibration model used for the run merely contained one effect ($N_{\text{AL}} = 1$ in Eq. 20), viz., the large-scale calibration $\Delta\eta$ in Eq. (15) accounting for geometric distortions of the CCDs and optical effects which are indistinguishable from geometric irregularities of the focal plane. On this account, we expect the simulated basic angle signal (see beginning of Sect. 7.2) to manifest itself through a corresponding time-dependence of the AL large-scale calibration parameters $\Delta\eta_{0fn0j}$. The asterisks in Fig. 15 show, for $f = \text{PFoV}$ and for each time interval j , the calibration parameter values of the demonstration solution averaged over the 62 CCDs (n). The solid line depicts the step-sinusoidal input basic-angle signal applicable for PFoV. As anticipated, the estimated calibration parameters are in very good agreement with




Fig. 15. Variation of the AL large-scale calibration parameters (averaged over all CCDs) in the preceding field of view (PFoV), as a function of the time since the beginning of the mission. The step-sinusoidal curve is the expected variation due to the simulated basic-angle variation having a period of 2.5 yr and an amplitude of $500 \mu\text{as}$, but constant within each 30 day interval. The asterisks show the results of the solution (one value per 30 day interval), and the circles show the differences on the magnified scale to the right. Thanks to the constraint in Eq. (16), the mean calibration parameters in the following field of view (FFoV) exactly mirror the displayed ones, and are therefore not shown.

the input signal: the RMS value of the differences is $0.20 \mu\text{as}$, corresponding to an RMS error of $0.40 \mu\text{as}$ for the basic angle offsets $\Delta\eta_j$ per calibration time interval (cf. Eq. 19). This is reasonably consistent with the expected precision of the large-scale calibration based on the total weight of the observations, as shown by the following calculation. The mean number of AL observations per calibration time interval and field of view is 1.33×10^7 (cf. Table 1). Assuming, as we did in Sect. 7.2.3, that an AL observation of average weight corresponds to a standard uncertainty of $\sigma_{\text{AL}} \approx 650 \mu\text{as}$, the expected precision of the basic angle determination is $2^{1/2} \times (650 \mu\text{as}) \times [1.33 \times 10^7]^{-1/2} \approx 0.25 \mu\text{as}$. The observed scatter, $0.40 \mu\text{as}$, is larger by roughly the same factor as found for the AL attitude errors.

The good agreement between the input calibration signal and the recovered parameters demonstrates the correct functioning of the generic calibration model (see Sect. 3.4) in this simple case. We expect that many more validation runs will be needed to confirm this result in more complex circumstances, i.e., with more calibration effects of different functional compositions and dependencies. A further important aspect is the practical study of possible hidden correlations and degeneracies of calibration with source, attitude, and global parameters which may not be fully obvious at the mathematical level.

7.2.5. Global results

Starting the iterations from a PPN γ value of 1.1, the purpose of running with the Global block was to see how such a grossly wrong initial value would affect the (overall) convergence rate and to what level the correct value could be recovered (recall that the input data were simulated using $\gamma = 1$).

Figure 16 presents the evolution of $g_0 = \gamma - 1$ during the run. In iteration 135 the parameter had settled on the value $g_0 =$




Fig. 16. Evolution of the estimated global parameter $g_0 = \gamma - 1$ (where γ is the PPN parameter) as a function of the iteration number. g_0 settles at a level of 2.6×10^{-5} from a starting value of $g_0 = 0.1$. The formal uncertainty of g_0 in this solution is 2.15×10^{-5} . Note that a linear scale is used for $|g_0| < 10^{-4}$ (the grey area), while a logarithmic scale is used outside of this interval.

2.62×10^{-5} , with a formal standard uncertainty²¹ of 2.15×10^{-5} . At $+1.2$ standard deviations, this value of g_0 is not significantly different from 0. The run shows that the system is capable of recovering the ‘correct’ value with an error compatible with the statistical uncertainty, which in this case is set by the photon noise of the individual observations.

The total astrometric weight of the AL observations in the demonstration run is $\approx 4000 \mu\text{as}^{-2}$. According to Hobbs et al. (2010) this should yield an RMS uncertainty in γ of about 3×10^{-5} , in reasonable agreement with the formal standard uncertainty given above and the parameter value obtained in the solution.

7.3. Processing times

The demonstration solution was run on an IBM cluster at ESAC using 14 nodes (out of 32 available), each node having two processors²² with four cores each; thus in total 112 CPUs were engaged. This configuration of 14 nodes is estimated to have a total floating point performance of $0.65 \text{ Tflop s}^{-1}$ (0.65×10^{12} floating point operations per second). One iteration took about 1 hr (with high CPU occupancy, typically 90%), so the total run time for 135 iterations was nearly 6 days, corresponding to about 3×10^{17} floating point operations (flop).

Scaled up to the projected 10^8 primary sources of the final AGIS run this would amount to 1.5×10^{19} flop. Using a more conservative estimate of 5×10^{19} flop to account for additional features not included in the demonstration run, this will require some 60 days on a typically targeted 10 Tflop s^{-1} machine. On the other hand, there could also be a significant saving in computing time due to the fact that the final solution will start off

from a previous solution, already very close to the final one; a smaller number of iterations might therefore suffice. In summary, the estimated processing time is clearly within the feasible range.

8. Conclusions

A fundamental part of the scientific data processing for the Gaia mission is the astrometric core solution, which will be run during and after the mission (ca. 2013–2020) with successively larger datasets and eventually encompassing at least some 100 million primary sources. This solution is central to the performance of the mission as a whole, since it not only provides the astrometric results for these primary sources, but also the reference frame (in the form of the instrument attitude as a function of time) and the geometric calibration of the instrument, for use by a large number of other processes in the overall scientific reduction of the Gaia data.

In order to accomplish the astrometric core solution, a software system known as AGIS (Astrometric Global Iterative Solution) is being built within Coordination Unit 3 (CU3) of the Gaia Data Processing and Analysis Consortium (DPAC). As detailed in this paper, the necessary mathematical models and numerical algorithms are well understood, and have been developed with sufficient rigour to allow the potential accuracy of Gaia to be fully exploited. Most critical parts of this system have been implemented, and numerous test runs have demonstrated the theoretical validity of the global iterative approach, as well as its practical feasibility in terms of data management and computations. While a number of additional features will have to be included in the software before it can be considered ready for the flight data, and many more complications will undoubtedly be discovered during the actual analysis of these data, all the fundamental parts of AGIS are already in place.

In 2011, with roughly two years left until the launch of the satellite and time to mature the concepts and software presented here into a robust operational system, we have no reason to doubt that AGIS will be able to compute an accurate astrometric solution, consistent with the ambitious goals of the Gaia mission.

Acknowledgements. The constant work of the Gaia Data Processing and Analysis Consortium (DPAC) has played an important part in this work. We are particularly indebted to CU2 for the production of independently simulated Gaia-like data for use in the system. The data simulations have been done in the supercomputer Mare Nostrum at Barcelona Supercomputing Center – Centro Nacional de Supercomputación (The Spanish National Supercomputing Center). Research and development in Sweden is kindly supported by the Swedish National Space Board (SNSB).

The successful development of concepts, algorithms and software incorporated in AGIS over a number of years, as well as the drafting of this manuscript, has benefited from the interaction with numerous colleagues, of which we especially wish to mention Sergei Klioner, Alex Bombrun, Alexey Butkevich, Jos de Bruijne, Berry Holl, Floor van Leeuwen, and François Mignard.

We thank the referee, Dr. F. van Leeuwen, for numerous constructive comments which helped to improve the original version of the manuscript.

Our special thanks go to Gaia’s former Project Scientist Michael Perryman, whose vision, leadership, and enthusiasm in the early years of the project laid the foundations for the excellent progress that is today seen throughout DPAC and with AGIS in particular.

References

- Axelsson, O. 1996, *Iterative Solution Methods* (Cambridge University Press)
- Bastian, U. 1995, in *ESA Special Publication*, Vol. 379, *Future Possibilities for astrometry in Space*, ed. M. A. C. Perryman & F. van Leeuwen, 99
- Bastian, U. & Biermann, M. 2005, *A&A*, 438, 745
- Björck, Å. 1996, *Numerical Methods for Least Squares Problems* (Society for Industrial and Applied Mathematics)

²¹ The formal uncertainty of g_0 was calculated as $N_{gg}^{-1/2}$ (Eq. 31) times the factor $(1 - \rho)^{-1/2} = \cot(\xi/2) = 2.414$, where $\xi = 45^\circ$ is the constant angle between Gaia’s spin axis (z) and the direction to the Sun; as explained by Hobbs et al. (2010), this factor takes into account the statistical correlation (ρ) between PPN γ and the parallaxes. The correction factor would not have been required, had the solution included the pseudo-parameter g_1 (Sect. 5.4).

²² Intel Nehalem EP Xeon (quad-core, 2.93 GHz), with 32 GB RAM

- Bombrun, A., Lindegren, L., Hobbs, D., et al. 2011, A&A, submitted
- Bombrun, A., Lindegren, L., Holl, B., & Jordan, S. 2010, A&A, 516, A77
- Chambers, J., James, D., Lambert, D., & Vander Wiel, S. 2006, *Statistical Science*, 21, 463
- de Boor, C. 2001, *A Practical Guide to Splines*, Rev. ed., Applied Mathematical Sciences, Vol. 27 (Springer)
- de Bruijne, J., Siddiqui, H., Lammers, U., et al. 2010, in IAU Symposium, Vol. 261, IAU Symposium, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, 331–333
- Dongarra, J., Bunch, J., Moler, C., & Stewart, G. 1979, *LINPACK Users' Guide* (Philadelphia, PA, USA: Society for Industrial and Applied Mathematics), 367
- Dravins, D., Lindegren, L., & Madsen, S. 1999, A&A, 348, 1040
- ESA. 1997, *The Hipparcos and Tycho Catalogues*, ESA SP-1200
- Feissel, M. & Mignard, F. 1998, A&A, 331, L33
- Gilbert, A. C., Kotidis, Y., Muthukrishnan, S., & Strass, M. J. 2002, in Proc. 28th International Conference on Very Large Data Bases, Hong Kong, China, 454
- Golub, G. H. & van Loan, C. F. 1996, *Matrix computations*, 3rd ed. (The Johns Hopkins University Press, Baltimore)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
- Greenbaum, A. 1997, *Iterative Methods for Solving Linear Systems* (Society for Industrial and Applied Mathematics)
- Greenwald, M. & Khanna, S. 2001, in Proc. ACM SIGMOD International Conference on Management of Data, Santa Barbara, CA, 58
- Gunn, J. E., Carr, M., Rockosi, C., et al. 1998, AJ, 116, 3040
- Gwinn, C. R., Eubanks, T. M., Pyne, T., Birkinshaw, M., & Matsakis, D. N. 1997, ApJ, 485, 87
- Hamilton, W. R. 1843, in *Proceedings of the Royal Irish Academy*, Vol. 2, 424
- Hansen, P. C. 1998, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, Monographs on Mathematical Modeling and Computation 4 (SIAM)
- Higham, N. 1990, in *Reliable Numerical Computation*, ed. M. G. Cox & S. J. Hammarling (Oxford University Press), 161
- Higham, N. J. 2002, *Accuracy and Stability of Numerical Algorithms*, 2nd edn. (Philadelphia, PA, USA: Society for Industrial and Applied Mathematics), 680
- Hobbs, D., Holl, B., Lindegren, L., et al. 2010, in IAU Symposium No. 261, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, 315
- Høg, E. 2008, in IAU Symposium No. 248, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, 300
- Høg, E., Fabricius, C., Makarov, V. V., et al. 2000, A&A, 355, L27
- Holl, B., Hobbs, D., & Lindegren, L. 2010, in IAU Symposium No. 261, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, 320
- Holl, B., Lindegren, L., & Hobbs, D. 2012, in preparation for A&A
- Huber, P. 1981, *Robust Statistics* (Wiley)
- Janesick, J. R. 2001, *Scientific charge-coupled devices* (SPIE Optical Engineering Press)
- Jordi, C., Gebran, M., Carrasco, J. M., et al. 2010, A&A, 523, A48+
- Kane, T. R., Likins, P. W., & Levinson, D. A. 1983, *Spacecraft dynamics* (McGraw Hill Book Company)
- Klioner, S. A. 2003, AJ, 125, 1580
- Klioner, S. A. 2004, Phys. Rev. D, 69, 124001
- Klumpp, A. R. 1976, J. Spacecraft, 13, 754
- Kopeikin, S. M. & Makarov, V. V. 2006, AJ, 131, 1471
- Laborie, A., Davancens, R., Pouney, P., et al. 2007, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6690, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- Lammers, U., Lindegren, L., O'Mullane, W., & Hobbs, D. 2009, in ASPC Series, Vol. 411, *Astronomical Data Analysis Software and Systems XVIII*, ed. D. A. Bohlender, D. Durand, & P. Dowler, 55
- Lawson, C. L. & Hanson, R. J. 1974, *Solving Least Squares Problems* (Prentice-Hall, New Jersey)
- Lindegren, L. 2005, in ESA Special Publication, Vol. 576, *The Three-Dimensional Universe with Gaia*, ed. C. Turon, K. S. O'Flaherty, & M. A. C. Perryman, 29
- Lindegren, L. 2010, in IAU Symposium No. 261, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, 296
- Lindegren, L., Babusiaux, C., Bailer-Jones, C., et al. 2008, in IAU Symposium No. 248, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, 217
- Lindegren, L. & Bastian, U. 2011, in EAS Publications Series, Vol. 45, 109–114
- Lindegren, L. & Dravins, D. 2003, A&A, 401, 1185
- Lindegren, L. & Kovalevsky, J. 1995, A&AS, 304, 189
- Lindegren, L. & Perryman, M. A. C. 1996, A&A, 116, 579
- Luri, X. & Babusiaux, C. 2011, in EAS Publications Series, Vol. 45, EAS Publications Series, 25–30
- Masana, E., Luri, X., Anglada-Escudé, G., & Llimona, P. 2005, in ESA Special Publication, Vol. 576, *The Three-Dimensional Universe with Gaia*, ed. C. Turon, K. S. O'Flaherty, & M. A. C. Perryman, 457–
- Mignard, F., Bailer-Jones, C., Bastian, U., et al. 2008, in IAU Symposium No. 248, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, 224
- Murray, C. A. 1983, *Vectorial astrometry* (Bristol: Adam Hilger)
- O'Mullane, W., Banday, A. J., Górski, K. M., Kunszt, P., & Szalay, A. S. 2001, in *Mining the Sky*, ed. A. J. Banday, S. Zaroubi, & M. Bartelmann, 638
- O'Mullane, W., Lammers, U., Bailer-Jones, C., et al. 2007, in *Astronomical Society of the Pacific Conference Series*, Vol. 376, *Astronomical Data Analysis Software and Systems XVI*, ed. R. A. Shaw, F. Hill, & D. J. Bell, 99
- O'Mullane, W., Lammers, U., Hernández, J., et al. 2010, in *Astronomical Society of the Pacific Conference Series*, Vol. 434, *Astronomical Data Analysis Software and Systems XIX*, ed. Y. Mizumoto, K.-I. Morita, & M. Ohishi, 135
- O'Mullane, W., Luri, X., Parsons, P., et al. 2011, *Experimental Astronomy*, 1, 10.1007/s10686-011-9241-6
- Perryman, M. A. C., de Boer, K. S., Gilmore, G., et al. 2001, A&A, 369, 339
- Press, W., Teukolsky, S., Vetterling, W., & Flannery, B. 2007, *Numerical Recipes: The Art of Scientific Computing*, 3rd edn. (Cambridge University Press)
- Prod'homme, T., Brown, A. G. A., Lindegren, L., Short, A. D. T., & Brown, S. W. 2011, MNRAS, 414, 2215
- Prod'Homme, T., Weiler, M., Brown, S. W., Short, A. D. T., & Brown, A. G. A. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7742, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Proust, D., Quintana, H., Carrasco, E. R., et al. 2006, A&A, 447, 133
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *Astronomy and Astrophysics*, 409, 523
- Schnabel, R. B. & Eskow, E. 1999, *SIAM Journal on Optimization*, 9, 1135
- Seabroke, G. M., Holland, A. D., Burt, D., & Robbins, M. S. 2009, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7439, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Short, A., Prod'homme, T., Weiler, M., Brown, S., & Brown, A. 2010, in *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, Vol. 7742, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Soffel, M., Klioner, S. A., Petit, G., et al. 2003, AJ, 126, 2687
- Stewart, G. 1998, *Matrix Algorithms: Basic decompositions* (Society for Industrial and Applied Mathematics)
- Turon, C., O'Flaherty, K. S., & Perryman, M. A. C., eds. 2005, *ESA Special Publication*, Vol. 576, *The Three-Dimensional Universe with Gaia*
- van der Vorst, H. 2003, *Iterative Krylov Methods for Large Linear Systems* (Cambridge University Press)
- van Leeuwen, F. 2005, A&A, 439, 805
- van Leeuwen, F. 2007, *Hipparcos, the New Reduction of the Raw Data*, *Astrophysics and Space Science Laboratory*, Vol. 350 (Springer)
- Wertz, J. R. 1978, *Spacecraft Attitude Determination and Control* (Kluwer Academic Publishers)

Appendix A: Quaternions

Quaternions form a non-commutative algebra in \mathbb{R}^4 . Invented by W. R. Hamilton in 1843 as an extension of the complex numbers (Hamilton 1843), their most common usage today is for representing spatial rotations in a particularly compact and convenient way, with applications for example in computer graphics and spacecraft attitude control. Quaternions can be introduced and understood in many different ways, with a correspondingly confusing multitude of notations and conventions. Here we give just a brief introduction to the subject, stating only the minimum results needed for our applications. For more details, see for example Wertz (1978) and Kane et al. (1983).

A.1. Quaternion algebra

A quaternion is a quadruple of real numbers for which the following algebraic operations are defined. For any quaternions $\mathbf{a} = \{a_x, a_y, a_z, a_w\}$ and $\mathbf{b} = \{b_x, b_y, b_z, b_w\}$ we have addition

$$\mathbf{a} + \mathbf{b} = \{a_x + b_x, a_y + b_y, a_z + b_z, a_w + b_w\}, \quad (\text{A.1})$$

multiplication

$$\mathbf{ab} = \left\{ \begin{aligned} &a_x b_w + a_y b_z - a_z b_y + a_w b_x, \\ &-a_x b_z + a_y b_w + a_z b_x + a_w b_y, \\ &a_x b_y - a_y b_x + a_z b_w + a_w b_z, \\ &-a_x b_x - a_y b_y - a_z b_z + a_w b_w \end{aligned} \right\}, \quad (\text{A.2})$$

and scalar multiplication

$$s\mathbf{a} = \{sa_x, sa_y, sa_z, sa_w\} \quad (\text{A.3})$$

for scalar s . Subtraction is analogous to addition, as derived from $\mathbf{a} - \mathbf{b} = \mathbf{a} + (-1)\mathbf{b}$. The conjugate of \mathbf{a} is

$$\mathbf{a}^* = \{-a_x, -a_y, -a_z, a_w\}, \quad (\text{A.4})$$

the norm (length) is

$$\|\mathbf{a}\| = (a_x^2 + a_y^2 + a_z^2 + a_w^2)^{1/2}, \quad (\text{A.5})$$

and the inverse (provided $\|\mathbf{a}\| > 0$) is

$$\mathbf{a}^{-1} = \|\mathbf{a}\|^{-2} \mathbf{a}^*. \quad (\text{A.6})$$

We have

$$(\mathbf{ab})^* = \mathbf{b}^* \mathbf{a}^*, \quad (\mathbf{ab})^{-1} = \mathbf{b}^{-1} \mathbf{a}^{-1}. \quad (\text{A.7})$$

Any non-zero quaternion can be normalized to unit length. In analogy with the notation for vector normalization, we use angular brackets for this operation:

$$\langle \mathbf{a} \rangle = \|\mathbf{a}\|^{-1} \mathbf{a}. \quad (\text{A.8})$$

The triplet of real numbers (a_x, a_y, a_z) can be regarded as the coordinates of the vector \mathbf{a} in some reference system $\mathbf{S} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$ (where $\mathbf{x}, \mathbf{y}, \mathbf{z}$ is a right-handed set of orthogonal unit vectors). Thus, we can write $\mathbf{a} = \{\mathbf{S}'\mathbf{a}, a_w\}$, where $\mathbf{S}'\mathbf{a} = [a_x \ a_y \ a_z]'$ constitutes the so-called vector part of the quaternion, and a_w the scalar part. Both scalars and vectors (in \mathbb{R}^3) can thus be seen as special cases of quaternions, namely, when the vector or the scalar part is zero. This allows us to write for example $\|\mathbf{a}\|^2 = \mathbf{aa}^* = \mathbf{a}^* \mathbf{a}$. In terms of the usual vector-scalar operations the quaternion multiplication can also be written as

$$\mathbf{ab} = \{\mathbf{S}'(\mathbf{a} \times \mathbf{b} + a_w \mathbf{b} + b_w \mathbf{a}), a_w b_w - \mathbf{a}' \mathbf{b}\}. \quad (\text{A.9})$$

Note that the vector part of a quaternion only makes sense when expressed in some coordinate system (\mathbf{S} in this example); a physical vector cannot be part of a quaternion.

A.2. Spatial rotations

Unit quaternions (of unit length) are convenient for representing orientations and spins of objects in three-dimensional space. This is compacter, numerically more stable and requires fewer arithmetic operations than the use of rotation matrices. Compared with the use of Euler angles, much fewer or no trigonometric functions need to be evaluated, and singularities are avoided.

According to Euler's rotation theorem any change in the orientation of an object can be described as a rotation by a certain angle around some fixed axis. Let this axis be represented by the unit vector \mathbf{u} and the rotation angle by ϕ , reckoned in the positive sense around the vector. In the given reference system

$\mathbf{S} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$, the rotation is then represented by the unit quaternion

$$\mathbf{q} = \left\{ \mathbf{S}'\mathbf{u} \sin \frac{\phi}{2}, \cos \frac{\phi}{2} \right\}. \quad (\text{A.10})$$

The useful property here is that a sequence of rotations is represented by the product of the corresponding unit quaternions (see Sect. A.3).

From Eq. (A.6) it follows that the inverse of a unit quaternion equals its conjugate, so

$$\mathbf{q}^{-1} = \mathbf{q}^* = \left\{ -\mathbf{S}'\mathbf{u} \sin \frac{\phi}{2}, \cos \frac{\phi}{2} \right\}, \quad (\text{A.11})$$

which represents a rotation by $-\phi$ around \mathbf{u} .

A rotation by the angle 2π around any axis is represented by the unit quaternion $\{0, 0, 0, -1\}$. Since this operation is physically equivalent to no rotation at all, it implies a sign ambiguity in the quaternion representation of any given rotation (modulo 2π). This is potentially a problem only when the quaternion is used to represent a continuously changing orientation as a function of time, as is the case for the Gaia attitude (Sect. 3.3). It is then necessary to ensure that no sign flips occur, e.g., when converting from some other representation of the orientations.

Equation (A.10) suggests an alternative, non-redundant, way of representing spatial rotation, namely by means of the components of the vector $\boldsymbol{\phi} = \mathbf{u}\phi$. For any continuous rotation the three components could be continuous functions of time. This formalism is mainly useful for small rotations ($\|\boldsymbol{\phi}\| \ll 1$); when applied for example to a spinning satellite the length of the vector may grow indefinitely, causing unacceptable numerical errors in finite arithmetic. For the arbitrary vector $\boldsymbol{\phi}$ we nevertheless introduce the special notation $\mathbf{Q}(\mathbf{S}'\boldsymbol{\phi})$ for the unit quaternion, in the reference system \mathbf{S} , representing a spatial rotation by the angle $\phi = \|\boldsymbol{\phi}\|$ about an axis parallel to $\boldsymbol{\phi}$:

$$\mathbf{Q}(\mathbf{S}'\boldsymbol{\phi}) = \begin{cases} \left\{ \mathbf{S}'\langle \boldsymbol{\phi} \rangle \sin \frac{\phi}{2}, \cos \frac{\phi}{2} \right\} & \text{if } \phi > 0, \\ \{0, 0, 0, 1\} & \text{if } \phi = 0. \end{cases} \quad (\text{A.12})$$

This notation is here only used when discussing the small rotational offset between the ICRS and AGIS frames in Sect. 6.1.1.

A.3. Vector and frame rotations

Vector rotation and frame rotation are not standard notions in vector or quaternion calculus, but we have found them helpful in order to clarify the relations between vectors and their representations in different reference systems.

Vector rotation. Let $\{\mathbf{S}'\mathbf{r}_0, 0\}$ be the quaternion representation in the reference system \mathbf{S} of the arbitrary vector \mathbf{r}_0 . Rotating the vector an angle ϕ around unit vector \mathbf{u} results in a new vector \mathbf{r}_1 , whose coordinates in \mathbf{S} can be calculated by two successive quaternion multiplications,

$$\{\mathbf{S}'\mathbf{r}_1, 0\} = \mathbf{q}\{\mathbf{S}'\mathbf{r}_0, 0\}\mathbf{q}^{-1}, \quad (\text{A.13})$$

where \mathbf{q} is given by Eq. (A.10) and $\mathbf{q}^{-1} = \mathbf{q}^*$. We call the operation in Eq. (A.13) for the *vector rotation* of \mathbf{r}_0 by the quaternion \mathbf{q} . This calculation requires the use of a particular reference system (\mathbf{S} in this example) in which to express the vectors and the quaternion; the resulting physical vector \mathbf{r}_1 is however independent of the reference system.

Applying n successive vector rotations, specified by the quaternions $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n$, gives the vector \mathbf{r}_n in

$$\{\mathbf{S}'_n \mathbf{r}_n, 0\} = \mathbf{q}_n \cdots \mathbf{q}_2 \mathbf{q}_1 \{\mathbf{S}'_0 \mathbf{r}_0, 0\} \mathbf{q}_1^{-1} \mathbf{q}_2^{-1} \cdots \mathbf{q}_n^{-1}. \quad (\text{A.14})$$

This is equivalent to a single vector rotation by $\mathbf{q} = \mathbf{q}_n \cdots \mathbf{q}_2 \mathbf{q}_1$. All the quaternions are here expressed in the fixed reference system \mathbf{S} , and the order of multiplication is from right to left.

Frame rotation. A more common application in astrometry is where the reference system itself is rotated, say from $\mathbf{S}_0 = [x_0 \ y_0 \ z_0]$ to $\mathbf{S}_1 = [x_1 \ y_1 \ z_1]$, by the quaternion \mathbf{q} . Given the coordinates $\mathbf{S}'_0 \mathbf{r}$ of the arbitrary vector \mathbf{r} in reference system \mathbf{S}_0 , we want to compute the coordinates $\mathbf{S}'_1 \mathbf{r}$ of the same physical vector in the rotated reference system. By applying a vector rotation to each of the basis vectors we find

$$\{\mathbf{S}'_1 \mathbf{r}, 0\} = \mathbf{q}^{-1} \{\mathbf{S}'_0 \mathbf{r}, 0\} \mathbf{q}, \quad (\text{A.15})$$

which we refer to as the *frame rotation* of \mathbf{r} by the quaternion \mathbf{q} .

It is important to note that the numerical representation of the quaternion \mathbf{q} , representing a frame rotation from \mathbf{S}_0 to \mathbf{S}_1 , is the same in the two frames. This follows because the components of the vector \mathbf{u} are invariant under a frame rotation about the vector itself, i.e., $\mathbf{S}'_1 \mathbf{u} = \mathbf{S}'_0 \mathbf{u}$. In Eq. (A.10) the vector part of \mathbf{q} can therefore be expressed in either of the two reference systems, i.e., we may take $\mathbf{S} = \mathbf{S}_0$ or $\mathbf{S} = \mathbf{S}_1$. The scalar part $\cos(\phi/2)$ is of course independent of the reference system.

Successive frame rotations can therefore be accomplished by referring each rotation to the current set of axes, which is usually precisely what is needed. Let $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n$ be a sequence of frame rotations, from \mathbf{S}_0 to \mathbf{S}_1 , and then from \mathbf{S}_1 to \mathbf{S}_2 , etc.; here \mathbf{q}_1 is expressed in \mathbf{S}_0 (or \mathbf{S}_1), \mathbf{q}_2 in \mathbf{S}_1 (or \mathbf{S}_2), and so on. The corresponding transformation of the coordinates of the vector \mathbf{r} is given by

$$\{\mathbf{S}'_n \mathbf{r}, 0\} = \mathbf{q}_n^{-1} \cdots \mathbf{q}_2^{-1} \mathbf{q}_1^{-1} \{\mathbf{S}'_0 \mathbf{r}, 0\} \mathbf{q}_1 \mathbf{q}_2 \cdots \mathbf{q}_n, \quad (\text{A.16})$$

equivalent to a single frame rotation by $\mathbf{q} = \mathbf{q}_1 \mathbf{q}_2 \cdots \mathbf{q}_n$. The quaternions are here expressed in the concurrent reference system, and the order of multiplication is from left to right.

A.4. Angular velocity

Let \mathbf{q} be a unit quaternion, expressed in the non-rotating reference system \mathbf{C} , and let us assume that \mathbf{q} is a differentiable function of time. The angular velocity $\boldsymbol{\Omega}$ associated with the time-dependent vector rotation by \mathbf{q} is the same for all vectors, and given by

$$\{\mathbf{C}' \boldsymbol{\Omega}, 0\} = 2 \dot{\mathbf{q}} \mathbf{q}^{-1}, \quad (\text{A.17})$$

where the dot signifies the time derivative. This result can be derived by taking the time derivative of the vector rotation formula for arbitrary vector \mathbf{r} , using $\dot{\mathbf{r}} = \boldsymbol{\Omega} \times \mathbf{r}$ and Eq. (A.9).

Let \mathbf{S} be the reference system obtained after rotation by \mathbf{q} . The coordinates of the angular velocity vector in the new system are found by performing a frame rotation of Eq. (A.17) by \mathbf{q} ; thus

$$\{\mathbf{S}' \boldsymbol{\Omega}, 0\} = \mathbf{q}^{-1} \{\mathbf{C}' \boldsymbol{\Omega}, 0\} \mathbf{q} = 2 \mathbf{q}^{-1} \dot{\mathbf{q}}. \quad (\text{A.18})$$

These expressions show, for example, how the instantaneous angular velocity of Gaia can be calculated either in the celestial reference system (CoMRS), using Eq. (A.17), or in the instrument system (SRS), using Eq. (A.18), from a knowledge of the attitude \mathbf{q} and its time derivative $\dot{\mathbf{q}}$.

Algorithm A.1 For given $\mathbf{A} = [A_{xx} \ A_{xy} \ A_{xz}; \ A_{yx} \ \cdots]$, this algorithm returns a unit quaternion $\mathbf{q} = \{q_x, q_y, q_z, q_w\}$ such that Eq. (A.19) is satisfied.

```

1:  $s \leftarrow A_{xx} + A_{yy} + A_{zz}$ 
2: for  $i = x, y, z$  do
3:    $|q_i| \leftarrow [(1 - s)/4 + A_{ii}/2]^{1/2}$ 
4: end for
5:  $q_w \leftarrow [(1 + s)/4]^{1/2}$ 
6: if  $|q_x| \geq \max(|q_y|, |q_z|)$  then
7:    $i \leftarrow x, j \leftarrow y, k \leftarrow z$ 
8: else if  $|q_y| \geq \max(|q_x|, |q_z|)$  then
9:    $i \leftarrow y, j \leftarrow z, k \leftarrow x$ 
10: else
11:    $i \leftarrow z, j \leftarrow x, k \leftarrow y$ 
12: end if
13:  $q_i \leftarrow |q_i| \text{sign}[A_{jk} - A_{kj}]$ 
14:  $q_j \leftarrow |q_j| \text{sign}[q_i(A_{ij} + A_{ji})]$ 
15:  $q_k \leftarrow |q_k| \text{sign}[q_i(A_{ik} + A_{ki})]$ 

```

A.5. The attitude matrix

For completeness, we give here the full relations between the attitude matrix defined in Sect. 3.3 and the quaternion representation of the attitude. Let $\mathbf{C} = [\mathbf{X} \ \mathbf{Y} \ \mathbf{Z}]$ be the celestial reference system (CoMRS; Sect. 3.1) and $\mathbf{S} = [\mathbf{x} \ \mathbf{y} \ \mathbf{z}]$ the instrument system (SRS). The attitude matrix \mathbf{A} is defined by Eq. (8). The attitude quaternion $\mathbf{q} = \{q_x, q_y, q_z, q_w\}$ gives the rotation from the CoMRS to the SRS, i.e., $\{\mathbf{C}' \mathbf{x}, 0\} = \mathbf{q} \{\mathbf{C}' \mathbf{X}, 0\} \mathbf{q}^{-1}$, etc. Then

$$\mathbf{A} = \begin{bmatrix} 1 - 2(q_y^2 + q_z^2) & 2(q_x q_y + q_z q_w) & 2(q_x q_z - q_y q_w) \\ 2(q_x q_y - q_z q_w) & 1 - 2(q_x^2 + q_z^2) & 2(q_y q_z + q_x q_w) \\ 2(q_x q_z + q_y q_w) & 2(q_y q_z - q_x q_w) & 1 - 2(q_x^2 + q_y^2) \end{bmatrix}. \quad (\text{A.19})$$

The conversion from \mathbf{A} to \mathbf{q} is less straightforward if we want to avoid numerical problems for certain attitudes. A stable algorithm was given by Klumpp (1976). In our notation, using pseudo-code, it is given by Algorithm A.1. Note that $\{q_x, q_y, q_z, q_w\}$ and $\{-q_x, -q_y, -q_z, -q_w\}$ correspond to the same \mathbf{A} , while the algorithm always returns a quaternion with $q_w \geq 0$; a reversal of the signs may therefore sometimes be required to ensure the temporal continuity of the quaternion components.

A.6. Differential rotation

Up until now, the quaternion formulae given in this Appendix hold rigorously for arbitrary rotations. We now derive a useful result, which however is only valid to first order in the (small) rotation angles. It can be used, for example, to compute the attitude error angles about the SRS axes, as was done in Sect. 7.2.3.

Let \mathbf{q}_0 and \mathbf{q}_1 be given unit quaternions, representing the two nearly co-aligned reference systems $\mathbf{S}_0 = [x_0 \ y_0 \ z_0]$ and $\mathbf{S}_1 = [x_1 \ y_1 \ z_1]$ with respect to some third (common) reference system such as the ICRS. It is required to express the difference between \mathbf{S}_1 and \mathbf{S}_0 by means of three small angles ϕ_x, ϕ_y, ϕ_z representing rotations about the axes in either system. More precisely, let $\boldsymbol{\phi} = [\phi_x \ \phi_y \ \phi_z]'$ be the spatial rotation that brings \mathbf{S}_0 into coincidence with \mathbf{S}_1 . Assuming that $\|\boldsymbol{\phi}\| \ll 1$ and neglecting terms of order $\|\boldsymbol{\phi}\|^2$, we have $\mathbf{S}_1 \simeq \mathbf{S}_0 + \boldsymbol{\phi} \times \mathbf{S}_0$, or

$$\mathbf{S}'_1 \mathbf{S}_0 \simeq \mathbf{I} + (\boldsymbol{\phi} \times \mathbf{S}_0)' \mathbf{S}_0 = \begin{bmatrix} 1 & \phi_z & -\phi_y \\ -\phi_z & 1 & \phi_x \\ \phi_y & -\phi_x & 1 \end{bmatrix}. \quad (\text{A.20})$$

According to Eq. (8) this matrix describes the orientation of \mathbf{S}_1 with respect to \mathbf{S}_0 . If \mathbf{d} is the quaternion representing a frame

rotation from \mathbf{S}_0 to \mathbf{S}_1 , we have (Sect. A.3) $\mathbf{q}_1 = \mathbf{q}_0 \mathbf{d}$ and hence

$$\mathbf{d} \equiv \{d_x, d_y, d_z, d_w\} = \mathbf{q}_0^{-1} \mathbf{q}_1. \quad (\text{A.21})$$

Given that $\|\phi\|$ is small, d_x, d_y and d_z are also small quantities, while $|d_w| \simeq 1$. Due to the sign ambiguity of the quaternion representation (Sect. A.2), d_w could however have either sign. Comparing Eqs. (A.19) and (A.20) we find, to first order in the small quantities,

$$\phi_x \simeq 2d_x d_w, \quad \phi_y \simeq 2d_y d_w, \quad \phi_z \simeq 2d_z d_w, \quad (\text{A.22})$$

where the factor d_w (being close to ± 1) guarantees that the angles obtain their correct signs. Equations (A.21)–(A.22) provide the required transformation from $(\mathbf{q}_0, \mathbf{q}_1)$ to (ϕ_x, ϕ_y, ϕ_z) .

Appendix B: Splines

A spline is a piecewise polynomial function $S(t)$ defined on some interval $[t_{\text{beg}}, t_{\text{end}}]$. That is, if the interval is divided into $K > 0$ sub-intervals by means of the instants t_k (called knots), such that $t_{\text{beg}} = t_0 < t_1 < \dots < t_K = t_{\text{end}}$, then in the sub-interval $t_k \leq t < t_{k+1}$ the spline function $S(t)$ equals some polynomial $P_k(t)$, $k = 0 \dots K-1$. The splines of interest here consist of polynomials of some fixed order M (or degree $M-1$); typically $M = 4$, corresponding to cubic splines. Moreover, the splines are usually maximally smooth, i.e., $S^{(m)}(t) \equiv d^m S/dt^m$ is continuous for $t_{\text{beg}} < t < t_{\text{end}}$ and $m = 0 \dots M-2$.

The K polynomials of order M require KM coefficients for their specification. For a maximally smooth spline, there are $M-1$ continuity conditions for each interior knot, namely $S^{(m)}(t_k^-) = S^{(m)}(t_k^+)$ for $m = 0 \dots M-2$ and $k = 1 \dots K-1$; thus a total of $(K-1)(M-1)$ constraints. The spline consequently has $KM - (K-1)(M-1) = K + M - 1$ degrees of freedom.²³ In the context of least-squares fitting, this equals the number of unknowns (parameters to fit), which we denote by N . In the following we take N and M to be the characteristic numbers of the spline, rather than K and M . For a maximally smooth spline we have $K = N - M + 1$.

B.1. B-splines

There are many different ways in which a spline could be represented (parametrized). The approach taken here is to consider $S(t)$ as a linear combination of some basis functions $B_n(t)$,

$$S(t) = \sum_{n=0}^{N-1} c_n B_n(t), \quad (\text{B.1})$$

where c_n are coefficients to be determined. Choosing the basis functions to have minimal support (see below) for the given order and smoothness leads to the so-called B-splines (de Boor 2001).

The B-splines are uniquely defined by the adopted knot sequence. In the following we shall use τ_n (rather than t_n) to denote

²³ When the spline is used for interpolation, it is typically chosen to go exactly through the $K+1$ values $S(t_k)$ for $k = 0 \dots K$. This leaves $(K+M-1) - (K+1) = M-2$ degrees of freedom. Thus, for a cubic spline ($M = 4$), two more conditions must be imposed for the spline to become unique. The most common choice is to make the second derivative vanish at t_{beg} and t_{end} . This defines the ‘natural’ interpolating spline. By contrast, when splines are least-squares fitted to data, as in the attitude determination problem, there are typically many data points per sub-interval, and no need for special endpoint conditions.




Fig. B.1. The first six B-splines of order $M = 4$ (cubic) defined on the regular knot sequence τ_0, τ_1, \dots . For better visibility, each B-spline is vertically displaced by one unit, and the non-zero parts are drawn in thick lines.

the knots associated with the B-splines. The reason is that the knot sequence for the B-splines has a slightly more general interpretation than just the simple division of $[t_{\text{beg}}, t_{\text{end}}]$ considered above. Moreover, when fitting the spline to given data points, this allows us to use t_k (for example) to denote the time associated with the k th data point, without ambiguity.

The knot sequence must be non-decreasing, $\tau_0 \leq \tau_1 \leq \tau_2 \leq \dots$, and at least two knots must be different. Very often we use a regular knot sequence in which $\tau_{n+1} = \tau_n + \Delta\tau$ for some $\Delta\tau > 0$ (the knot interval). Figure B.1 shows the first six cubic B-splines defined on a regular knot sequence. Note that the non-zero part of each B-spline, shown in heavy line, stretches over M consecutive knot intervals. We use the convention that the B-splines are labelled with the same index as the left-most knot of its non-zero part; therefore, the support of $B_n(t)$ is (τ_n, τ_{n+M}) .

It is not possible to construct a maximally smooth spline function of order M that is non-zero over less than M consecutive knot intervals. This is what we mean by the statement that B-splines have minimal support: they could not be shortened without losing some smoothness. This is an important property for the least-squares fitting, as shown by the following consideration. Least-squares fitting of (B.1) to a given set of data points (t_k, z_k) (where $t_k \in [t_{\text{beg}}, t_{\text{end}}]$ for each k) will be done by forming normal equations. The normal equations matrix N is a symmetric, positive definite matrix of dimension²⁴ $N \times N$. Since N may be very large, it is desirable that the matrix is sparse, i.e., that most elements are zero. It is easy to see that element N_{ij} will be non-zero as soon as $B_i(t_k)B_j(t_k) \neq 0$ for some data point k . To make the matrix maximally sparse, we should therefore choose basis functions with minimal support. B-splines have minimal support and are therefore ideal for least-squares fitting using sparse matrix algebra.

Since the support of a B-spline of order M extends over at most M consecutive knot intervals, we have $N_{ij} = 0$ for $|i -$

²⁴ In the attitude determination problem, each of the four components of the quaternion is represented by a spline, so the number of parameters is actually $4N$ and the normal matrix is of dimension $4N \times 4N$. Alternatively, we may take the elements of N to be sub-matrices of dimension 4×4 . The indices i and j used below refer to these sub-matrices.

Algorithm B.1 For given spline order M , knot sequence $\{\tau_n\}$, time t , and left index ℓ (such that $\tau_\ell \leq t < \tau_{\ell+1}$), this algorithm returns $\{b_0 \dots b_{M-1}\}$ such that $b_0 = B_{\ell-M+1}(t)$, \dots , $b_{M-1} = B_\ell(t)$.

```

1:  $b_0 \leftarrow 1$ 
2: for  $i = 0$  to  $M - 2$  do
3:    $R_i \leftarrow \tau_{\ell+i+1} - t$ 
4:    $L_i \leftarrow t - \tau_{\ell-i}$ 
5:    $s \leftarrow 0$ 
6:   for  $j = 0$  to  $i$  do
7:      $u \leftarrow b_j / (R_j + L_{i-j})$ 
8:      $b_j \leftarrow s + R_j \times u$ 
9:      $s \leftarrow L_{i-j} \times u$ 
10:  end for
11:   $b_{i+1} \leftarrow s$ 
12: end for

```

$j| > M - 1$. This shows that N is a symmetric banded matrix with (half-) bandwidth equal to $M - 1$. Cholesky decomposition preserves the band structure of the matrix and is therefore ideal for solving the normals; it is also numerically very efficient.

B.2. Calculation of B-splines

At any given point t there are at most M non-zero B-splines, namely $B_{\ell-M+1}(t)$, $B_{\ell-M+2}(t)$, \dots , $B_\ell(t)$, where ℓ is the 'left index' of t , such that $\tau_\ell \leq t < \tau_{\ell+1}$. They can be computed simultaneously in a numerically stable way by means of de Boor's algorithm (de Boor 2001), which is given as pseudo-code in Algorithm B.1. If needed, the derivatives of the B-splines with respect to t can be computed simultaneously with little additional effort.

By inspection of the algorithm it is found that the knots used for computing the B-spline values in the interval $[\tau_\ell, \tau_{\ell+1})$ are $\tau_{\ell-M+2}$ through $\tau_{\ell+M-1}$. For example, with reference to Fig. B.1 (with $M = 4$), the B-splines between τ_3 and τ_4 (i.e., for left index $\ell = 3$) depend on $\{\tau_1, \tau_2, \dots, \tau_6\}$, but not on τ_0 or τ_7 , even though $B_0(t)$ in general depends on τ_0 and $B_3(t)$ depends on τ_7 .

B.3. Use of multiple knots

Algorithm B.1 works also in the case when several consecutive knots are placed at the same t coordinate. Having a knot of multiplicity m (i.e., m knots at the same t) removes the continuity for derivatives of order $M - m$ and higher. The normal situation is that the knots have multiplicity 1, which means that the spline is continuous at the knots up to and including its $(M-2)$ th derivative, but discontinuous in its $(M-1)$ th derivative. By inserting multiple knots at some specific instant, one allows the spline to become less smooth at this point. For example, in a cubic spline ($M = 4$) a triple knot ($m = 3$) allows the first derivative to become discontinuous at that point, while leaving the spline function itself continuous. In the Gaia attitude processing, multiple knots will be used for modelling the effects of micrometeoroid impacts, which cause (almost) instantaneous changes in the angular velocity, corresponding to discontinuities in the first derivative of the attitude spline.

Multiple knots may also be used at the endpoints of the spline interval $[t_{\text{beg}}, t_{\text{end}}]$. At any point in this interval, there must be exactly M non-zero B-splines in order that a linear combination of them should be able to represent an arbitrary spline of order M . Again, with reference to Fig. B.1 (for $M = 4$), we see that this is the case to the right of τ_3 (or τ_{M-1} in general). Thus we should put $\tau_{M-1} = t_{\text{beg}}$. The first $M - 1$ knots can in principle




Fig. B.2. Illustration of the knot placement for a spline of order M (e.g., $M = 4$ for a cubic spline) with N degrees of freedom. $[t_{\text{beg}}, t_{\text{end}}]$ is the spline interval over which the spline is fitted to given data points. The end knots τ_{M-1} and τ_N are at the endpoints of the spline interval. The $N - M$ interior knots are chosen to give the spline the desired flexibility, including multiple knots where required. The placement of the anterior ($\tau \leq t_{\text{beg}}$) and posterior knots ($\tau \geq t_{\text{end}}$) is in principle arbitrary: it does not change the fitted spline in $[t_{\text{beg}}, t_{\text{end}}]$, but may affect the condition number of the least-squares fit. The parameters of the fitted spline are the coefficients c_0, \dots, c_{N-1} of the B-splines $B_0(t)$ through $B_{N-1}(t)$.

be placed anywhere, as long as $\tau_0 \leq \tau_1 \leq \dots \leq \tau_{M-1}$: any such arrangement will produce M non-zero B-splines in the next sub-interval (τ_{M-1}, τ_M) , and although the B-splines are different depending on the arrangement of the knots, they are always linearly independent and therefore can be used as a basis for the spline. In particular, it is possible to put the first M knots at the same point, i.e., $\tau_0 = \tau_1 = \dots = \tau_{M-1} = t_{\text{beg}}$.

Corresponding considerations apply to the right limit of the spline interval: with N degrees of freedom, the support of the last B-spline $B_{N-1}(t)$ extends from τ_{N-1} to τ_{N+M-1} . The spline interval must end at $\tau_N = t_{\text{end}}$, and the remaining $M - 1$ knots can be placed anywhere provided that $\tau_N \leq \tau_{N+1} \leq \dots \leq \tau_{N+M-1}$. In particular, it is possible to have an M -fold knot at the endpoint of the spline interval, i.e., $t_{\text{end}} = \tau_N = \tau_{N+1} = \dots = \tau_{N+M-1}$. Figure B.2 summarizes the placement of knots in relation to a given spline interval for given order M and degrees of freedom N (number of spline coefficients).

Although the placement of the first and last $M - 1$ knots is arbitrary, and does not change the resulting spline between t_{beg} and t_{end} , their placement does affect the numerical stability of the resulting least-squares system. Butkevich & Klioner (unpublished technical note) has pointed out that collapsing the anterior and posterior knots into the end knots, so that $\tau_0 = \tau_1 = \dots = \tau_{M-1} = t_{\text{beg}}$ and $t_{\text{end}} = \tau_N = \tau_{N+1} = \dots = \tau_{N+M-1}$ results in a system with much smaller condition number than a regular sequence extending beyond the endpoints. For a cubic spline ($M = 4$) the spline interval then begins and ends with 4-fold knots as illustrated in Fig. B.3. However, the use of data segmentation, as described in Sect. 5.2.2, may not permit this device.

Appendix C: A robust Cholesky algorithm for positive semidefinite matrices without pivoting

C.1. The use of normal equations

The least-squares problems considered in this paper are solved by the method of normal equations (here denoted $N\mathbf{x} = \mathbf{b}$), using the Cholesky algorithm to decompose the symmetric normal matrix N . This method is known to be computationally efficient




Fig. B.3. Definition of a regular knot sequence for fitting a cubic spline (order $M = 4$) in the interval $[t_{\text{beg}}, t_{\text{end}}]$. The spline interval is divided into $N - 3$ knot intervals of equal length $\Delta\tau = (t_{\text{end}} - t_{\text{beg}})/(N - 3)$.

and accurate for well-conditioned problems (e.g., Stewart 1998), and is moreover well adapted to in-place manipulation of sparse matrices such as the band matrix obtained when fitting B-splines (Fig. 5)

The method of normal equations for the general least-squares problem is usually discouraged in the literature, due to the much superior stability and accuracy of alternative methods operating directly on the observation equations $\mathbf{Ax} \simeq \mathbf{h}$ (where $\mathbf{N} = \mathbf{A}'\mathbf{A}$ and $\mathbf{b} = \mathbf{A}'\mathbf{h}$), e.g., using Householder orthogonal transformations (e.g., Björck 1996). However, when working with very large problems that are inherently well-posed, thanks to a good design of experiment and reduction model, our experience is that the method of normal equations is nearly always adequate in terms of accuracy, and then has the edge over other methods in terms of speed, storage and simplicity of the code. Moreover, in these problems, iterative improvement of the solution is usually required for other reasons (non-linearity, elimination of outliers), which partly compensates for the loss of precision when forming the normal equations.

C.2. The Cholesky algorithm

The standard Cholesky algorithm (e.g., Björck 1996; Golub & van Loan 1996) requires that \mathbf{N} is positive definite, which is always the case for a well-conditioned least-squares problem. Given \mathbf{N} and \mathbf{b} , the solution of the system $\mathbf{Nx} = \mathbf{b}$ proceeds in three steps:

- C1. Use the Cholesky algorithm to find the unique upper-diagonal matrix \mathbf{U} , with positive diagonal entries, such that $\mathbf{N} = \mathbf{U}'\mathbf{U}$.
- C2. Solve the lower-triangular system $\mathbf{U}'\mathbf{y} = \mathbf{b}$.
- C3. Solve the upper-triangular system $\mathbf{U}\mathbf{x} = \mathbf{y}$.

The matrix \mathbf{U} (or its transpose) is known as the Cholesky factor or square root of \mathbf{N} . As all the computations can be made in-place, we have symbolically

$$[\mathbf{N} | \mathbf{b}] \xrightarrow{\text{C1}} [\mathbf{U} | \mathbf{b}] \xrightarrow{\text{C2}} [\mathbf{U} | \mathbf{y}] \xrightarrow{\text{C3}} [\mathbf{U} | \mathbf{x}]. \quad (\text{C.1})$$

The extension to an arbitrary number of right-hand sides (to the right of the vertical line) is trivial. For example, the inverse \mathbf{N}^{-1} can be obtained by inserting the identity matrix for \mathbf{b} . We also note that C1 and C2 are mathematically equivalent to pre-multiplying the matrix and right-hand side, respectively, by $(\mathbf{U}')^{-1}$. Performing C2 on the unit vector $\mathbf{e}_i = [0, 0, \dots, 1, \dots, 0]'$ (with 1 in position i) thus produces $\tilde{\mathbf{e}}_i = (\mathbf{U}')^{-1}\mathbf{e}_i$ such that $\tilde{\mathbf{e}}_i'\tilde{\mathbf{e}}_j = \mathbf{e}_i'\mathbf{U}^{-1}(\mathbf{U}')^{-1}\mathbf{e}_j = \mathbf{e}_i'\mathbf{N}^{-1}\mathbf{e}_j = (\mathbf{N}^{-1})_{ij}$. Selected elements or sub-matrices of \mathbf{N}^{-1} can thus be obtained without computing the full matrix.

Algorithm C.1 Returns upper-triangular \mathbf{U} such that $\mathbf{U}'\mathbf{U} = \mathbf{N}$, where $\mathbf{N} \in \mathbb{R}^{n \times n}$ is symmetric and positive semidefinite. Also computes an estimate d of the rank defect of \mathbf{N} .

```

1:  $\mathbf{U} \leftarrow \mathbf{N}$ ,  $d \leftarrow 0$ 
2: for  $j = 0$  to  $n - 1$  do
3:   for  $i = 0$  to  $j$  do
4:      $s \leftarrow \sum_{k=0}^{i-1} U_{ki}U_{kj}$ 
5:     if  $i < j$  then
6:       if  $U_{ii} > 0$  then
7:          $U_{ij} \leftarrow (U_{ij} - s)/U_{ii}$ 
8:       else
9:          $U_{ij} \leftarrow 0$ 
10:      end if
11:    else
12:      if  $U_{jj} - s > 0$  then
13:         $U_{jj} \leftarrow \sqrt{U_{jj} - s}$ 
14:      else
15:         $U_{jj} \leftarrow 0$ ,  $d \leftarrow d + 1$ 
16:      end if
17:    end if
18:  end for
19:  for  $i = j + 1$  to  $n - 1$  do
20:     $U_{ij} \leftarrow 0$ 
21:  end for
22: end for

```

Algorithm C.2 Returns \mathbf{y} such that $\mathbf{U}'\mathbf{y} = \mathbf{b}$, where $\mathbf{U} \in \mathbb{R}^{n \times n}$ is upper-triangular and $\mathbf{b} \in \mathbb{R}^n$.

```

1:  $\mathbf{y} \leftarrow \mathbf{b}$ 
2: for  $i = 0$  to  $n - 1$  do
3:   if  $U_{ii} > 0$  then
4:      $y_i \leftarrow (y_i - \sum_{k=0}^{i-1} U_{ki}y_k)/U_{ii}$ 
5:   else
6:      $y_i \leftarrow 0$ 
7:   end if
8: end for

```

The above three steps are accomplished by Algorithms C.1–C.3 for arbitrary symmetric positive definite \mathbf{N} . (Actually, these algorithms include the non-standard modification discussed below in order to handle semidefinite matrices gracefully.) A few remarks should be made concerning its practical implementation. First, the matrices \mathbf{U} and \mathbf{N} in C1, \mathbf{y} and \mathbf{b} in C2, and \mathbf{x} and \mathbf{y} in C3 can share the same memory if it is acceptable that \mathbf{U} overwrites \mathbf{N} , etc (in-place calculation). Second, since \mathbf{N} is symmetric, only the upper-diagonal part of it (N_{ij} for $i \leq j$) is used in C1, and similarly for \mathbf{U} . For large systems one can therefore save roughly half the memory by storing only the upper-diagonal parts of \mathbf{N} and \mathbf{U} , e.g., as one-dimensional arrays. The code in lines 19–21 of Algorithm C.1 is then irrelevant. Third, if \mathbf{N} is a ‘skyline matrix’ with envelope E_j (i.e., $N_{ij} = 0$ for $i < E_j$) then \mathbf{U} has the same envelope: the Cholesky decomposition gives no fill-in above the envelope. This allows to store and decompose certain sparse matrices very efficiently, such as the band matrix in Fig. 5.

C.3. Application to semidefinite systems

In several of our applications the normal matrix is however not positive definite, either from a lack of observations (e.g., data gaps in the attitude spline representation) or by design (e.g., for the calibration parameters). Application of the standard Cholesky algorithm in such cases results in an exception which may be non-trivial to handle (for example, by changing

Algorithm C.3 Returns \mathbf{x} such that $\mathbf{U}\mathbf{x} = \mathbf{y}$, where $\mathbf{U} \in \mathbb{R}^{n \times n}$ is upper-triangular and $\mathbf{y} \in \mathbb{R}^n$.

```

1:  $\mathbf{x} \leftarrow \mathbf{y}$ 
2: for  $i = n - 1$  to 0 do
3:   if  $U_{ii} > 0$  then
4:      $x_i \leftarrow (x_i - \sum_{k=i+1}^{n-1} U_{ik}x_k)/U_{ii}$ 
5:   else
6:      $x_i \leftarrow 0$ 
7:   end if
8: end for

```

the knot sequence of the attitude spline) and which may leave the partially solved system in an undefined state. However, that the Cholesky algorithm fails does not mean that there is no solution to the normal equations: on the contrary, there is an infinitude of solutions and the problem is rather which one to pick. Indeed, in many situations we could in principle accept *any* valid solution; for example, when the null space of the problem is known a priori, and we are prepared to handle the associated non-uniqueness of the solution (cf. Sect. 6.1). For these and several other reasons it is advantageous if the computation can be continued in some sensible way, while of course noting the detected rank deficiency.

A number of methods are available to handle rank-deficient or ill-conditioned least-squares problems. Singular Value Decomposition (SVD; e.g., Björck 1996; Golub & van Loan 1996; Press et al. 2007) is the method often recommended; it provides the unique ‘pseudo-solution’ with the smallest Euclidean norm. However, SVD is computationally expensive and the pseudo-solution is not necessarily better than any other valid solution to the singular least-squares problem.

By construction, the normal matrix $\mathbf{N} = \mathbf{A}'\mathbf{A}$ is positive semidefinite: $\mathbf{x}'\mathbf{N}\mathbf{x} = \|\mathbf{A}\mathbf{x}\|^2 \geq 0$ for any \mathbf{x} (cf. footnote 7). A modification of the standard Cholesky algorithm allows the decomposition in C1 to be made also for such matrices, although \mathbf{U} is no longer unique; similarly C2 and C3 can be modified to produce a valid (if non-unique) solution to the normal equations $\mathbf{N}\mathbf{x} = \mathbf{b}$. For example, the LINPACK routine xCHDC (Dongarra et al. 1979) implements a robust version of the Cholesky algorithm, using complete pivoting (i.e., a simultaneous permutation of the rows and columns of \mathbf{N}) to generate the unique square root with non-zero elements only in the first r rows, if $r < n$ is the rank of the matrix (for a detailed analysis, see Higham 2002). Other modifications of the Cholesky algorithm (e.g., Schnabel & Eskow 1999) also uses pivoting.

Permuting the rows and/or columns of the matrix is highly undesirable in the present applications. For example, when applied to a band matrix (such as Fig. 5) it likely destroys the band structure, and in general prevents the envelope-based storing of the sparse matrices \mathbf{N} and \mathbf{U} outlined above. While pivoting is never needed for the Cholesky factorization of a positive definite matrix, it is thought to be an essential ingredient in modified algorithms aimed at more general symmetric matrices. A simple modification of the Cholesky algorithm, which makes it applicable to the semidefinite case without pivoting was described by Lawson & Hanson (1974, Eq. 19.5); see also Golub & van Loan (1996, Eq. 4.2.11), who however warn that “it may be preferable to incorporate pivoting”. Nevertheless we have implemented the corresponding modifications in Algorithms C.1–C.3 without pivoting, and find that they work quite well in our applications. Algorithm C.1 includes a rough estimation of the rank defect $d = n - r$.

The numerical accuracy of the decomposition in Algorithm C.1 was tested in MATLAB for a range of rank-deficient random positive semidefinite matrices \mathbf{N} (patterned after Higham 1990) by computing the quantity $\rho_N = \|\mathbf{N} - \mathbf{U}'\mathbf{U}\|_F / (u\|\mathbf{N}\|_F)$ as a measure of the relative error in units of the floating point precision. Here $\|\mathbf{N}\|_F = [\text{trace}(\mathbf{N}'\mathbf{N})]^{1/2}$ is the Frobenius norm of \mathbf{N} , and $u = 2^{-52}$ is the unit roundoff (Golub & van Loan 1996) of the double precision floating point arithmetic used. We find that the present algorithm, without pivoting, performs almost as well as LINPACK’s xCHDC with complete pivoting, as judged from the statistics of our ρ_N compared with the corresponding quantity ρ_k reported by Higham. However, C.1 is much less useful as a rank-revealing algorithm – the estimated rank defect is often much too small.

Similarly, in order to check the numerical validity of the solution to the rank-deficient normal equations computed by C.1–C.3, we made the following experiments, using the same matrices as above. For random vectors \mathbf{x} we first computed $\mathbf{b} = \mathbf{N}\mathbf{x}$, then used C.1–C.3 to recover a solution $\tilde{\mathbf{x}}$ (usually very different from \mathbf{x}). Finally we computed $\rho_b = \|\mathbf{b} - \mathbf{N}\tilde{\mathbf{x}}\|_F / (u\|\mathbf{b}\|_F)$. We find that our algorithm performs almost as well as MATLAB’s basic solution $\mathbf{xtilde} = \mathbf{N} \backslash \mathbf{b}$, and much better than the minimum norm solution $\mathbf{xtilde} = \text{pinv}(\mathbf{N}) * \mathbf{b}$ (with default tolerance). For example, the 99th percentile of ρ_b was $\sim 10^4$ for C.1–C.3, $\sim 10^3$ for MATLAB’s backslash (\backslash) operator, and $\sim 10^{11}$ (!) when using pinv .

We conclude from these limited experiments that the present version of the Cholesky algorithm is a useful, simple and efficient substitute for much more sophisticated algorithms applicable in the semidefinite case. Since it does not use pivoting, it preserves the envelope of the matrix and is therefore especially suited for banded matrices and envelope-based sparse matrix methods. It provides an estimate of the rank of the matrix, which however is rather unreliable. In the positive definite case the algorithm is equivalent to the standard Cholesky method.

Appendix D: Complexities beyond the basic modelling

Section 3 describes a set of baseline models for the sources, attitude, and geometrical instrument, which are believed to be realistic enough to serve as an acceptable first-order approximation of the actual data for primary sources. Due to the many complexities of the real satellite and its operation, as well as the physical environment in space, there are however many additional effects that may affect the astrometric results at the μas level, and which have to be considered in the final modelling. In this Appendix we discuss some of the known effects that will be addressed in future versions of the astrometric solution.

D.1. Chromaticity

Although the Gaia telescopes are all-reflective, with no refractive elements in the optical paths to the astrometric field, they are nevertheless not completely achromatic. In the presence of odd wavefront errors, such as coma, the centroids of the optical images do in fact depend on the wavelength, and hence on the spectral energy distributions of the observed sources.²⁵ For the typical wavefront errors expected in the astrometric field of

²⁵ ‘Centroid’ should here not be understood as the centre of gravity of the optical image; rather, it is a non-trivial function of the light distribution, similar to the estimation of the image location κ in Sect. 3.5.

Gaia (about 50 nm RMS), the AL centroid shift from an early-type star to a late spectral type could amount to several mas. This systematic effect, known as chromaticity, can therefore be many times larger than the photon-statistical uncertainty of the estimated image location (cf. Table 1). It is thus essential to have a very good calibration of the chromaticity, for which it is necessary to know the spectral energy distribution of every observed source. This is obtained from the photometric observations in the BP and RP fields (see Fig. 3).

Chromaticity is eliminated in the CCD signal analysis by using a Line Spread Function (LSF), $L(x)$ in Eq. (23), which is correspondingly shifted depending on the (known) spectrum of the source. The resulting AL pixel coordinate κ is therefore in principle achromatic, and the effect need not be further considered in the astrometric solution. However, as mentioned in Sect. 3.4, it is envisaged to have diagnostic colour-dependent terms in the geometric instrument model of the astrometric solution. These calibration parameters should obtain negligible values in the solution if the chromaticity has been correctly accounted for in the calibration of $L(x)$. Conversely, non-zero values can be used to improve the LSF calibration in the next processing cycle.

D.2. Charge Transfer Inefficiency of the CCDs

The CCD signal model in Eq. (23) assumes a perfectly linear detector, which is not exactly the case for the real detectors and especially not in the presence of radiation damage on the CCDs. Traps in the silicon substrate, produced by particle radiation in the space environment, will capture charges during the TDI operation of the CCDs, and release them with delays ranging from a fraction of the TDI period to minutes. The charge capture and release processes introduce a number of phenomena, collectively referred to as Charge Transfer Inefficiency (CTI = 1 – CTE, where CTE is the Charge Transfer Efficiency; Janesick 2001). CTI will affect all kinds of observations in Gaia (astrometric, photometric, spectroscopic). The most important phenomena for the astrometric observations are the apparent charge loss (because part of the charges are released outside the observed window) and centroid shift (because some charges are released with a delay of one or a few TDI periods). These and more general effects of the radiation damage are the subject of extensive theoretical and experimental studies within the Gaia community (e.g., Seabroke et al. 2009; Prod’homme et al. 2010; Prod’homme et al. 2011). The adopted method to handle these effects in the Gaia data processing is by means of forward modelling using a so-called Charge Distortion Model (CDM; Short et al. 2010). In the context of the CCD signal model of Sect. 3.5, the CDM may be represented by the (non-linear) operator D :

$$\lambda_k \equiv E(N_k) = D \left[\{\lambda_{k'}^0\}_{k' \leq k} \mid \Psi \right], \quad (\text{D.1})$$

where λ_k^0 is the signal model at sample k in the absence of radiation-damage effects, e.g., according to Eq. (23). The variable Ψ represents the state of the CCD, e.g., in terms of how much radiation damage it has suffered. Note that the expected value of sample k depends on the CCD illumination history up to and including sample k , which is expressed by the CDM taking as input the (undamaged) value not only for sample k but also for the preceding samples ($k' \leq k$).²⁶ One of the methods

employed for mitigating CTI effects in Gaia is through a periodic (e.g., once per second) electronic injection of charges in a few consecutive TDI lines. As the lines of charges travel across the CCD, most of the harmful traps are temporarily filled, thus reducing the CTI of subsequent charge transfers (Laborie et al. 2007). The method has the additional benefit of periodically resetting the illumination history of the pixels, so that in Eq. (D.1) only the samples since the previous charge injection need to be considered (Short et al. 2010).

The CDM depends on a moderate number of parameters that will be estimated in parallel with the LSF (or PSF) calibration prior to the astrometric solution (the ‘Instrument response parameters’ in Fig. 1). *In principle*, the subsequently estimated image location κ should then not only be achromatic, as discussed in Appendix D.1, but also free of CTI effects, so that the astrometric solution can use a purely geometric instrument model, as required by the primary source model. Although this is obviously a highly idealised condition, it is nevertheless what the final data analysis must aim to achieve.

For the simple image of a primary source, the centroid shift due to the CTI depends mainly on the magnitude of the source, the time since the previous charge injection, and the accumulated radiation dose experienced by the CCD (Prod’homme et al. 2011). It is expected that imperfections in the CDM calibration will likewise show up in systematic shifts depending primarily on these (known) quantities, and can be represented by a set of diagnostic calibration functions in the generic calibration model of Sect. 3.4. Non-negligible values of the diagnostic parameters in the astrometric solution indicate that the CDM is not doing its job properly, and they can then be used to improve the model. The reader is referred to the cited papers for quantitative information on the expected level of CTI effects in Gaia data, the effectiveness of different mitigation strategies, and the associated performance degradations.

D.3. Effects of the finite CCD integration time

Up until now we have regarded the astrometric observations of Gaia as instantaneous measurements of the crossings of the source images over the fiducial ‘observation line’ (Fig. 4) at the centre of the CCD (or of the gated portion of the CCD) in the AL direction. In reality, due to the finite integration time (T) of the CCD observations, any measurement clocked into the CCD readout register at time t actually depends on the average attitude and source position over the preceding integration time interval, $[t - T, t]$. The time delay is, to first order, taken into account by associating the measurement with the observation time $t - T/2$ (Sect. 3.5). Following Bastian & Biermann (2005) we should, more precisely, assume that the observed location κ of the image centre in the pixel stream is a weighted mean of the *instantaneous* location $\kappa_*(t)$ of the optical image relative to the charge image during the preceding integration interval:

$$\kappa = \frac{\int_0^T e(\tau) \kappa_*(t - \tau) d\tau}{\int_0^T e(\tau) d\tau}, \quad (\text{D.2})$$

where $e(\tau)$ is the (nominally flat) ‘exposure function’ for look-back time τ , i.e., the rate at which electrons are produced, and

For the sake of illustration, the centroid could for example be the point obtained by fitting a Gaussian PSF to the image.

²⁶ A further complication is that most AF observations are binned in the AC direction before read-out, while the CTI effects operate indepen-

dently on each pixel column. Thus the CDM should ideally be applied to the two-dimensional charge image, and the distorted charge image then binned for comparison with the one-dimensional data. This requires the use of a PSF replacing the LSF in Eq. (23).

transported to the read-out register, for constant illumination. The instantaneous location is given by

$$\kappa_*(t) = s[\eta_*(t) - \eta_{fng}] + k(t), \quad (\text{D.3})$$

where s is the local scale factor (pixels per radian), $\eta_*(t)$ the instantaneous AL field angle of the optical image centre, and η_{fng} the AL coordinate of the fiducial observation line for the appropriate AC coordinate, field of view, etc. The function $k(t)$ is the inverse of t_k , relating the time coordinate to the TDI period index k . For the present discussion $k(t)$ is regarded as a continuous function, ignoring the step-wise transportation of the charge image in TDI mode.²⁷

Recalling that $\eta_*(t)$ is decreasing with time (cf. Fig. 3), while $k(t)$ is increasing, it is seen that $\kappa_*(t)$ remains approximately constant throughout the integration. Let us denote by t_c the exact time when the centre of the optical image crosses the fiducial observation line, so that $\eta_*(t_c) = \eta_{fng}$, and let $\kappa_c = k(t_c)$ be the corresponding pixel coordinate. If the speed of the optical image exactly matches the speed of the charge image, or $s\dot{\eta}_* + \dot{k} = 0$, it is seen that $\kappa_*(t)$ is indeed constant and equal to κ_c . Let us now consider what happens if there is a mismatch between the speeds. This could be caused by (i) a deviation in the local scale value s due to optical distortion; (ii) a non-nominal local scan rate; and (iii) that the object itself has significant motion (e.g., an asteroid). Assuming that \dot{k} is constant and adopting a Taylor series expansion for the AL field angle over the exposure time, we have

$$k(t) = \kappa_c + (t - t_c)\dot{k}, \quad (\text{D.4})$$

$$\eta_*(t) = \eta_{fng} + (t - t_c)\dot{\eta}_* + \frac{1}{2}(t - t_c)^2\ddot{\eta}_* + \dots \quad (\text{D.5})$$

Inserting in Eq. (D.3), and assuming that s is constant across the section of the CCD in question, we obtain by means of Eq. (D.2)

$$\kappa = \kappa_c + (s\dot{\eta}_* + \dot{k})e_1 + \frac{1}{2}s\ddot{\eta}_*e_2 + \dots, \quad (\text{D.6})$$

where

$$e_m = \frac{\int_0^T e(\tau)(T/2 - \tau)^m d\tau}{\int_0^T e(\tau) d\tau}, \quad m = 1, 2, \dots \quad (\text{D.7})$$

are the normalized moments of $e(\tau)$ relative to the exposure mid-time at $\tau = T/2$. For a constant and matching image speed we recover $\kappa = \kappa_c$ as expected. In the general case of imperfect speed matching and non-constant scan rate there is a difference which should be taken into account in the astrometric solution. If we assume that s is known, the speed mismatch $s\dot{\eta}_* + \dot{k}$ can be computed for every observation, and the factor e_1 can then be estimated as an instrument calibration parameter using the generic calibration model in Sect. 3.4. e_1 will depend (at least) on the CCD and gate used; but due to the accumulating radiation damage it is likely to evolve with time and could possibly have a magnitude-dependent component as well. In the next (quadratic) term we may know $\ddot{\eta}_*$ (from the attitude) and $e_2 \approx T^2/12$ (for constant exposure function) to sufficient accuracy that it might

be applied as a correction to the observed κ . However, since most observations are ungated (giving maximum T), it may be better to adopt the uncorrected κ for this maximum T as defining the *effective* attitude,²⁸ and only correct the gated observations for the difference in e_2 ; hence they, too, will refer to the effective attitude. A complication is that $\ddot{\eta}_*$ in Eq. (D.6) should be computed from (unknown) *physical* attitude, but to first order it can be obtained from the effective attitude. Attitude irregularities on time scales shorter than T add further complications (see Appendix D.4), but it is unlikely that higher-order terms in Eq. (D.6) can profitably be accounted for.

D.4. Attitude irregularities

The basic attitude model described in Sect. 3.3 uses a spline representation which is (normally) continuous in the angles specifying the instantaneous orientation of the instrument, as well as in the first $M - 2$ derivatives of the angles, where M is the order of the spline (typically cubic splines are used, for which $M = 4$). The actual (physical) attitude is much more complex and in particular there may be discontinuities and irregularities on time scales that are too short to be adequately represented by a spline with the knot separations considered in the basic model. Low-frequency perturbations ($\lesssim 0.01$ Hz) are of no concern here, as they can be perfectly represented by splines. The most important contributors to perturbations at higher frequencies are thruster noise in the micro-propulsion system used for the attitude control; the discrete and partially stochastic nature of the control system (for example that the commanded thrusts are updated once per second); micrometeoroid impacts on the satellite; and various dynamical effects such as fuel sloshing and structural vibrations.

The high-frequency perturbations due to the thruster noise and control system generate angular jitter of the physical attitude that has a significant amplitude relative to the astrometric accuracies ultimately aimed at, but still small in comparison with the AL pixel size (≈ 59 mas). Thanks to the TDI integration this high-frequency jitter is largely removed from the effective attitude by the averaging over the exposure time T . As a result, the shortest knot interval needed to accurately represent the effective attitude is also of the order of the exposure time, or about 5 s. The optimum knot interval may be longer, depending on the number and magnitudes of the primary sources available for the attitude determination, and on the actual level of perturbations.

The expected frequency of micrometeoroid impacts of various sizes can be predicted from the known velocity and mass spectrum of interplanetary particles. Each impact produces a quasi-instantaneous change in the angular velocity of the satellite, while the attitude angles are continuous across the impact time. It is expected that a few hundred impacts will occur every year producing a change in the AL angular velocity exceeding 1 mas s^{-1} (which should be easily detectable), with the frequency roughly inversely proportional to the minimum change considered. Discontinuities in the *physical* attitude rate can be represented in the spline model by inserting multiple knots at the estimated times of impact, t_i (Sect. 5.2.6). However, the *effective* attitude will instead see a linear change in the attitude rates over an interval equal to the exposure time T , centred on t_i , which requires that multiple knots are inserted both at $t_i - T/2$ and $t_i + T/2$. (This treatment becomes more problematic in connection with

²⁷ The corresponding expression in Bastian & Biermann (2005) is their Eq. (6), in which $k(t)$ is the (integer) index of the last TDI clock stroke before time t . Thus their $\kappa_*(t)$ oscillates with an amplitude of ± 0.5 unit for every TDI period. The continuous approximation adopted in our Eq. (D.3) is acceptable since we are always considering integrals covering a whole number of oscillations.

²⁸ The effective attitude is then the physical attitude convolved with the (average) exposure function for maximum T . It corresponds closely to the ‘astrometric attitude’ introduced by Bastian & Biermann (2005).

gated observations, when T is non-nominal.) Impacts will be detected by inspecting the observation residuals in connection with the attitude updating (Sect. 5.2.5).

The detailed re-examination of the Hipparcos attitude by van Leeuwen (2005, 2007) revealed numerous discontinuities of the along-scan attitude angle (scan phase) of several tens of mas. A large fraction of them could be linked to the beginning or end of eclipses experienced by Hipparcos in its highly elliptic orbit around the Earth. A likely cause is thermal re-adjustment of the solar-panel hinges, following a sudden change in temperature. As there is no change in the net angular momentum, but only a re-distribution of inertia, the attitude rates are the same before and after a discontinuity. For Gaia it is estimated that such ‘clanks’ will be negligible, but the attitude processing should nevertheless be capable of identifying instances, should they occur, and to take appropriate measures. Due to the finite CCD integration time, the apparent effects of a clank will be two discontinuities in the attitude rates, equal but of opposite sign, and separated by the integration time T . Again, this can be handled by suitable modification of the knot sequence of the attitude spline. Like the micrometeoroid impacts, clanks will be detected during the attitude updating by means of the characteristic patterns that they generate in the observation residuals versus time.

Appendix E: Tables of acronyms and variables

Table E.1 is a list of acronyms used in the paper. Table E.2 lists the most important variables, with a short description and a reference to where it is introduced or explained.

Table E.1. List of acronyms

Acronym	Description
AC	ACross scan direction (Fig. 3)
ACF	ACross scan coordinate in the Following FoV
ACP	ACross scan coordinate in the Preceding FoV
AF	Astrometric Field
AGIS	Astrometric Global Iterative Solution
AGN	Active Galactic Nucleus
AL	ALong scan direction (Fig. 3)
ASI	Accelerated Simple Iteration
BAM	Basic-Angle Monitor (Fig. 3)
BCRS	Barycentric Celestial Reference System
BP	Blue Photometer (Fig. 3)
CCD	Charge-Coupled Device
CDM	Charge Distortion Model
CFS	Calibration Faint Star
CG	Conjugate Gradient
CPU	Central Processing Unit
CoMRS	Centre of Mass Reference System
CTE	Charge Transfer Efficiency (of a CCD)
CTI	Charge Transfer Inefficiency (of a CCD)
CU2	DPAC Coordination Unit 2, ‘Data Simulations’
CU3	DPAC Coordination Unit 3, ‘Core Processing’
DPAC	Data Processing and Analysis Consortium
EADS	European Aeronautic Defence and Space company
ESA	European Space Agency
ESAC	European Space Astronomy Centre
FFoV	Following Field of View (Fig. 2)
FPA	Focal Plane Assembly (Fig. 3)
FoV	Field of View
GCRS	Geocentric Celestial Reference System
HEALPix	Hierarchical Equal-Area iso-Latitude Pixelisation
IAU	International Astronomical Union
ICRS	International Celestial Reference System
LSF	Line Spread Function
NSL	Nominal Scanning Law
PFoV	Preceding Field of View (Fig. 2)
PPN	Parametrised Post-Newtonian (relativity formalism)
PSF	Point Spread Function
RMS	Root-Mean-Square
RP	Red Photometer
RSE	Robust Scatter Estimate (footnote 18)
RVS	Radial Velocity Spectrometer (Fig. 3)
SI	Simple Iteration
SM	Sky Mapper
SRS	Scanning Reference System (Fig. 2)
SVD	Singular Value Decomposition
TB	TeraByte
TCB	Barycentric Coordinate Time
TDI	Time-Delayed Integration (CCD operation mode)
VLBI	Very Long Baseline Interferometry
WFS	WaveFront Sensor (Fig. 3)
XML	eXtensible Markup Language

Table E.2. List of mathematical variables

Var.	Description	Ref.
A_u	the astronomical unit	Sect. 3.2
\mathbf{A}	attitude matrix	Eq. (8)
\mathbf{a}	galactocentric acceleration vector	Sect. 6.1.4
\mathbf{a}	attitude parameters	Sect. 3.3
\mathbf{a}_n	quaternion coefficient for $B_n(t)$	Eq. (10)
$B_n(t)$	B-spline function	Sect. B.1
\mathbf{B}	iteration matrix	Sect. 4.5
\mathbf{b}	right-hand side of normal equations	Eq. (30)
$\mathbf{b}_G(t)$	barycentric coordinate of Gaia at time t	Eq. (4)
\mathbf{C}	celestial reference system	Sect. 3.1
\mathbf{C}	constraint matrix for the calibration parameters	Eq. (22)
\mathbf{c}	calibration parameters	Sect. 3.4
\mathbf{D}	non-linear Charge Distortion Model	Eq. (D.1)
\mathbf{d}	update vector	Sect. 4.5
\mathbf{E}	expectation operator	Eq. (23)
$e(t)$	CCD exposure function	Appendix D.3
\mathbf{e}	error vector	Sect. 4.5
f	field index (± 1 for preceding/following)	Eq. (13)
\mathbf{f}	detector coordinates	Eq. (1)
$\mathbf{f}_P, \mathbf{f}_F$	preceding, following viewing direction	Fig. 2
G	Gaia broadband magnitude	Sect. 3.4
g	gate index	Eq. (13)
\mathbf{g}	global parameters	Sect. 3.4
\mathbf{h}	auxiliary data, e.g., ephemerides	Eq. (7)
\mathbf{I}	the identity matrix	Sect. 4.5
i	subscript denoting some source	Sect. 3.2
j	‘short’ calibration time interval index	Eq. (15)
K	no. of TDI periods for CCD integration	Sect. 3.5
\mathbf{K}	preconditioner matrix	Sect. 4.5
k	TDI index in CCD pixel stream	Sect. 3.5
k	‘long’ calibration time interval index	Eq. (15)
k	iteration index	Eq. (40)
L	Line Spread Function	Eq. (23)
L_r^*	shifted Legendre polynomial of degree r	Eq. (15)
l	subscript denoting some observation	Sect. 3.6
ℓ	left index in a knot sequence	Sect. B.2
M	spline order – $M = 4$ for cubic spline	App. B
N	number of degrees of freedom for a spline	App. B
\mathbf{N}	normal equations matrix	Eq. (30)
n	attitude parameter index	Eq. (10)
n	CCD index	Eq. (14)
n	dimension of the normal matrix \mathbf{N}	Eq. (35)
\mathbf{n}	nuisance parameters	Eq. (1)
\mathbf{p}_i	normal-triad component	Eq. (5)
Q	objective function to be minimized	Eq. (24)
\mathbf{q}	attitude quaternion	Eq. (10)
\mathbf{q}_i	normal-triad component	Eq. (5)
r	degree of the small-scale calibration polynomial	Eq. (15)
\mathbf{r}_i	normal-triad component	Eq. (5)
R_l	residual of observation l	Sect. 3.6
\mathbf{S}	Scanning Reference System	Sect. 3.1
s	local scale factor	Eq. (D.3)
\mathbf{s}	astrometric (‘source’) parameters	Sect. 3.2
\mathbf{s}_i	the astrometric parameters for source i	Sect. 3.2
T	light integration (exposure) time on CCD	Sect. 3.4, D.3
t_k	time for sample k	Sect. 3.5
t_l	time for observation l	Sect. 3.5
t_{ep}	reference epoch for astrometric parameters	Sect. 3.2
t_{fr}	reference epoch for frame rotator	Sect. 6.1
t_p	reference epoch for non-rotating ICRS sources	Sect. 6.1.4

Table E.2. List of mathematical variables (continued)

Var.	Description	Ref.
t_{PM}	reference epoch for moving ICRS sources	Sect. 6.1.4
U	relegation factor for primary-star selection	Sect. 6.2.2
\mathbf{u}_i	proper direction to source i	Sect. 3.2
$\bar{\mathbf{u}}_i$	geometric direction to source i	Sect. 3.2
\mathbf{V}	nullspace of the normal matrix	Sect. 4.5
\mathbf{W}_i	weight matrix for source i	Eq. (57)
w_l	downweighting factor for observation l	Sect. 3.6, 5.1.2
W_l	statistical weight of observation l	Sect. 3.6
\mathbf{x}	differential parameter vector	Sect. 4.3
$[X\ Y\ Z]$	celestial reference system (ICRS or CoMRS)	Sect. 3.1
$[x\ y\ z]$	Scanning Reference System (SRS)	Fig. 2
α	flux (image parameter)	Eq. (23)
α_i	right ascension of source i at t_{ep}	Sect. 3.2
β	background level (image parameter)	Eq. (23)
Γ_{c}	(conventional) basic angle	Eq. (13)
γ	PPN curvature parameter	Eq. (88)
δ_i	declination of source i at t_{ep}	Sect. 3.2
ϵ_a	excess attitude noise	Eq. (28)
ϵ_i	excess source noise (for source i)	Eq. (28)
ϵ_l	excess noise in observation l	Eq. (28)
$\boldsymbol{\varepsilon}$	orientation correction (frame rotator)	Sect. 6.1
ζ	across-scan (AC) field angle	Eq. (12)
η	along-scan (AL) field angle	Eq. (13)
κ	along-scan (AL) pixel coordinate (image location)	Eq. (23)
λ	regularization parameter for the attitude update	Eq. (81)
λ_k	intensity for CCD sample values	Eq. (23)
μ	across-scan (AC) pixel coordinate (image location)	Sect. 3.4
$\mu_{\alpha*i}$	proper motion in α ($\times \cos \delta_i$) for source i	Sect. 3.2
$\mu_{\delta i}$	proper motion in δ for source i	Sect. 3.2
μ_{ri}	radial proper motion for source i	Sect. 3.2
$\boldsymbol{\mu}_0$	proper motion due to galactocentric acceleration	Sect. 6.1.4
ν	number of degrees of freedom	Sect. 5.1.2
ϖ_i	parallax for source i	Sect. 3.2
ϱ	normalized RSE error of astrometric parameters	Sect. 7.2.2
σ_l	formal standard uncertainty for observation l	Sect. 3.6
τ_n	attitude spline knot n	Sect. B.1
φ	along-scan instrument angle	Eq. (12)
ω	spin correction (frame rotator)	Sect. 6.1